UNIVERSIDAD POLITÉCNICA DE MADRID

ESCUELA TÉCNICA SUPERIOR
DE INGENIEROS DE TELECOMUNICACIÓN

TECHNOLOGICAL CONTRIBUTIONS TO IMAGING
RADARS IN THE MILLIMETER-WAVE BAND

TESIS DOCTORAL

Federico Antonio García Rial

INGENIERO DE TELECOMUNICACIÓN

2019

DEPARTAMENTO DE SISTEMAS, SEÑALES Y RADIOCOMUNICACIONES

ESCUELA TÉCNICA SUPERIOR DE INGENIEROS DE TELECOMUNICACIÓN
UNIVERSIDAD POLITÉCNICA DE MADRID

TECHNOLOGICAL CONTRIBUTIONS TO IMAGING RADARS IN THE MILLIMETER-WAVE BAND

TESIS DOCTORAL

Autor:
Federico Antonio García Rial
Ingeniero de Telecomunicación
Universidad Politécnica de Madrid

Director:
Jesús Grajal de la Fuente
Catedrático del Dpto. de Señales, Sistemas y Radiocomunicaciones
Universidad Politécnica de Madrid

2019

TESIS DOCTORAL

**TECHNOLOGICAL CONTRIBUTIONS TO IMAGING RADARS IN THE
MILLIMETER-WAVE BAND**

AUTOR: Federico Antonio García Rial

DIRECTOR: Jesús Grajal de la Fuente

Tribunal nombrado por el Mgfc. y Excmo. Sr. Rector de la Universidad Politécnica de Madrid, el día de de 2019.

PRESIDENTE:.....

SECRETARIO:.....

VOCAL:.....

VOCAL:.....

VOCAL:.....

SUPLENTE:.....

SUPLENTE:.....

Realizado el acto de defensa y lectura de la Tesis el día de de 2019, en Madrid.

Calificación:

EL PRESIDENTE

LOS VOCALES

EL SECRETARIO

Abstract

Recent research in the millimeter- and submillimeter-wave frequency bands has driven the maturity of its technologies to a level where integrated and commercial radiofrequency components are now readily available. Due to this development the fundamental applications of these frequency bands (like spectroscopy, radioastronomy, and wireless communication) have benefited from increasing performance, at a fraction of the cost. This is also the case for imaging applications, such as concealed weapon detection, which is progressively becoming a viable alternative to traditional security screenings. Nonetheless, some technological progress is still needed for security imaging in this frequency range to fully reach its potential, specially regarding critical and diverse areas such as component cost, power consumption, imaging speed, or privacy concerns.

The main goal of this Ph.D. thesis can be found within this framework, and it is to develop those technological improvements and solutions necessary to enhance the performance of security imaging radars in the millimeter- and submillimeter-wave bands. This objective has been accomplished in two particular radar systems: an active imager at 300 GHz (developed previously to the start of this thesis), and a hybrid system using active and passive sensors at 122 GHz and 94 GHz, respectively. Throughout the development, a special emphasis is placed on achieving the desired results with the use of commercially-available components, in an effort to truly provide a viable alternative. Although both systems are meant for security screening purposes, their radiofrequency and processing architectures present various differences according to their operational requirements. In addition to their security-related function, both systems have also been employed in other applications, such as vital sign monitoring.

The thesis has been entirely carried out at the *Grupo de Microondas y Radar* of the *Universidad Politécnica de Madrid*, and is divided into two main parts. The first half deals with the improvements made to the 300 GHz imaging system, including a real-time signal and image processing architecture and the initial tests of a multistatic configuration. The second half is focused on the development of the hybrid imager at 122 GHz and 94 GHz, detailing individual sensor characterizations, simultaneous performance, and an image fusion technique to present the captured data. In each section, pending improvement actions are also studied and described for future research and work.

Keywords:

Concealed weapon detection, imaging radar, millimeter-wave, submillimeter-wave, terahertz.

Resumen

Investigaciones recientes en la banda de frecuencias de ondas milimétricas y submilimétricas han impulsado la madurez de sus tecnologías hasta un nivel donde componentes de radiofrecuencia comerciales e integrados ya están disponibles. Debido a este desarrollo las aplicaciones fundamentales de estas bandas de frecuencia (como la espectroscopia, la radioastronomía, y las comunicaciones inalámbricas) se han beneficiado de un rendimiento en incremento, por una fracción del coste habitual. Este es también el caso de las aplicaciones de imagen, como la detección de amenazas ocultas, que está progresivamente convirtiéndose en una alternativa viable a las inspecciones de seguridad tradicionales. Sin embargo, aún es necesario cierto progreso tecnológico para que las aplicaciones de seguridad basadas en imágenes en este rango de frecuencias alcancen su potencial completo, especialmente en áreas tan críticas y diversas como el coste de componentes, el consumo, la velocidad de refresco de imágenes, o las preocupaciones por privacidad.

La meta principal de esta tesis doctoral se puede encontrar dentro de este marco, y es desarrollar las mejoras y soluciones tecnológicas necesarias para elevar el rendimiento de los radares de imagen para seguridad en las bandas de milimétricas y submilimétricas. Este objetivo se ha cumplido para dos sistemas radar en particular: un sistema activo a 300 GHz (desarrollado previamente a esta tesis), y un sistema híbrido usando sensores activos y pasivos a 122 GHz y 94 GHz, respectivamente. Durante el desarrollo, se ha puesto un énfasis especial en alcanzar los resultados deseados mediante el uso de componentes comercialmente disponibles, en un esfuerzo por verdaderamente proporcionar una alternativa viable. Aunque ambos sistemas están pensados para las inspecciones de seguridad, sus arquitecturas de radiofrecuencia y procesamiento presentan varias diferencias según sus requisitos operacionales. Además de en su función de seguridad, ambos sistemas se han probado en otras aplicaciones, como la detección remota de señales vitales.

Esta tesis se ha llevado a cabo completamente en el Grupo de Microondas y Radar de la Universidad Politécnica de Madrid, y se divide en dos partes principales. La primera mitad trata sobre las mejoras realizadas al sistema de imagen a 300 GHz, incluyendo una arquitectura de procesamiento de señal e imagen en tiempo real y las pruebas iniciales de una configuración multiestática. La segunda mitad se centra en el desarrollo del sistema de imagen híbrido a 122 GHz y 94 GHz, detallando las caracterizaciones individuales de los sensores, su rendimiento simultáneo, y una técnica de fusión de imágenes para presentar los datos capturados. En cada sección, se estudian y detallan acciones de mejora pendientes para futuro trabajo e investigación.

Palabras Clave:

Detección de amenazas ocultas, radar de imagen, ondas milimétricas, ondas submilimétricas, terahercios.

Acknowledgements

Quisiera expresar mi agradecimiento a todas las personas que a través de su ayuda han hecho posible esta tesis. En primer lugar, debo agradecer a mi tutor, D. Jesús Grajal de la Fuente, su apoyo durante la duración del trabajo aquí presentado y el haberme dado la oportunidad de llevar a cabo esta tesis.

Extiendo ese agradecimiento a todo el Grupo de Microondas y Radar de la Universidad Politécnica de Madrid. En especial a mis compañeros durante estos años: Dani, Luis, Carlos, Gorka, Mario, Rodrigo, Alejandro, Luis Alberto, Marta, Clara y Carlos, en quienes me he apoyado durante este tiempo.

También debo mi gratitud al Ministerio de Educación, Cultura y Deporte, por la concesión de la Beca de Formación de Profesorado Universitario, cuya aportación económica me ha permitido realizar esta tesis. También agradezco su ayuda a la Dirección General de Tráfico, que a través de su programa de ayuda a proyectos de investigación financió mi primer año como doctorando, y a NVIDIA, ya que gracias a su *research grant* obtuvimos nuevos recursos de computación. Además, debo agradecer a todos los proyectos nacionales e internacionales que han costado las diversas publicaciones y conferencias realizadas durante esta tesis.

Como no podía ser de otra manera, un agradecimiento especial merece la paciencia, el ánimo, y la comprensión recibidos de mis padres y mi hermana, quienes han vivido esta tesis casi tanto como yo.

Y en especialmente gracias a Ana, para quien todas las páginas de esta tesis quedarían cortas en agradecimiento.

Contents

1	Introduction	1
1.1	Overview of the Millimeter-wave Spectrum	1
1.1.1	History of the Millimeter-wave Spectrum	1
1.1.2	Phenomenology of the Millimeter-wave Spectrum	2
1.1.3	Applications in the Millimeter-wave Spectrum	5
1.1.4	Outlook and Trends	7
1.2	Imaging in the Millimeter-wave Spectrum	8
1.2.1	Comparison between Active and Passive Imaging	8
1.2.2	Imager Architecture	12
1.2.3	State-of-the-art Imaging Systems	20
1.3	Thesis Motivation and Objectives	22
1.4	Publications and Author's Contributions	23
1.5	Research Projects Involved	24
1.6	Thesis Structure	25
2	300 GHz Single-sensor Security Imager	27
2.1	Introduction to a 300 GHz Standoff Imaging Radar	27
2.1.1	System Description	28
2.2	GPU-based Real-time Image Processing	33
2.2.1	Initial Processing Architecture	36
2.2.1.1	GPU-code Profiling	43
2.2.1.2	CPU-code Profiling	48
2.2.1.3	Complete Profiling	51
2.2.2	Improved Processing Architecture	51
2.2.2.1	GPU-code Profiling	53
2.2.2.2	CPU-code Profiling	56
2.2.2.3	Complete Profiling	57
2.2.3	System Performance	59
2.3	Millimeter-wave Scattering and Multistatic Configuration	64
2.3.1	Millimeter-wave Scattering Behavior	67
2.3.2	Imaging Simulations	78
2.3.3	Evaluation of Multistatic Imaging for Security Screening	82
2.4	Future Work	85
2.4.1	Real-Time Imaging	85
2.4.1.1	GPU-based 3-D Reconstruction	85
2.4.1.2	Multi-GPU Processing Architecture	86
2.4.1.3	Hardware Optimizations	88
2.4.2	Multistatic Scattering and Imaging	89
3	122 GHz and 94 GHz Multi-sensor Security Imager	91
3.1	122 GHz Radar Characterization and Performance	91
3.1.1	122 GHz MMIC-based Radar	91
3.1.2	Evaluation of Radar Performance	96

3.1.2.1	SNR	97
3.1.2.2	Resolution	101
3.1.2.3	Ranging Accuracy and Sensor Stability	104
3.1.3	122 GHz Waveforms Considered	113
3.2	Development of a Hybrid Imaging System	114
3.2.1	Initial Imager and Passive Receiver Characterization	116
3.2.1.1	Temperature Resolution	119
3.2.1.2	Image SNR and Contrast	121
3.2.1.3	Cross-Range Resolution	131
3.2.1.4	Spot Size	131
3.2.2	Characterization of the Hybrid Imager's Active Sensor	132
3.2.2.1	Cross-Range Resolution	134
3.2.2.2	Spot Size	136
3.2.3	Sensor Compatibility for Simultaneous Scanning	136
3.2.4	Image Fusion	141
3.2.4.1	Additional Active Sensor Image Processing	144
3.2.4.2	Additional Passive Sensor Image Processing	148
3.2.4.3	Imaging Scenarios	148
3.2.5	Array Scanning	150
3.3	Future Work	154
4	Automatic Threat Detection (ATD)	157
4.1	ATD in Active Millimeter-wave Images	157
4.1.1	Image Dataset	158
4.1.2	Threat Detection with a Convolutional Neural Network (CNN)	162
4.1.3	Detection Results	170
4.2	ATD in Passive Infrared Images	182
4.2.1	Image Dataset and Network Architecture	182
4.2.2	Detection Results	186
4.3	Future Work	193
4.3.1	ATD in Active Images	194
4.3.2	ATD in Passive Images	195
5	Other Applications	197
5.1	Data Acquisition on Plastic Guides	197
5.1.1	Radar Performance	204
5.1.2	Guide Measurements	210
5.2	Vital Signs Monitoring	216
5.2.1	For the 122 GHz radar	219
5.2.2	For the 300 GHz radar	223
5.3	Future Work	225
5.3.1	Data Acquisition	225
5.3.2	Vital Sign Monitoring	227
5.3.3	Future Applications for Millimeter-wave Radars	229
5.3.3.1	Vibration Detection	229
5.3.3.2	Multi-node Sensor Network	232
6	Conclusions and Future Work	235
6.1	Conclusions	235
6.2	Future Work	236
	Bibliography	237
	Annexes	265

A	Reference Data for a 300 GHz Imaging Radar	267
A.1	Image Representation Considerations	267
A.2	Automatic Target Recognition and Clustering	274
A.3	Image Database and Quality Metric	279
A.4	Minor Changes to the Original System	281
B	RF Components for a 300 GHz Imaging Radar	285
B.1	x6 Active Frequency Multiplier	285
B.2	Bandpass Filter	292
C	Synchronization Electronics for a 300 GHz Imaging Radar	295
C.1	Scanning Electronics	295
C.2	Bandpass Sampling	305
D	Modifications to a 122 GHz Commercial Radar	309
E	Equipment and Material for Data Acquisition	317
F	Material Analysis with FMCW Radars	325
F.1	Analysis of Absorption in Plastics	325
F.2	Analysis of Absorption in Common Fabrics	333
F.3	Future Work	336
G	Bandwidth and Size Improvements for the 300 GHz Imaging Radar	339
	Acronyms	345

List of Figures

1.1	Region of the electromagnetic spectrum which includes the mm-wave, submm-wave, and THz bands.	1
1.2	Absorption coefficient of deionized water according to the signal's frequency, reproduced from [16].	3
1.3	Atmospheric attenuation throughout the spectrum, with marked absorption peaks from oxygen and water molecules, reproduced from [19]. Above ~ 300 GHz windows become less broad as closer absorption lines appear and baseline attenuation increases.	3
1.4	Absorption of various fabrics throughout the THz spectrum, modified from [23].	4
1.5	(a) Active imaging scenario, where received power is dominated by reflections of the transmitted signal. (b) Passive imaging scenario, where received power presents approximately equally-powerful contributions from background, clutter, and the subject's own emissions.	10
1.6	(a) Passive outdoors imaging scenario, where received temperature from threat and subject presents a good contrast thanks to the sky's cold radiation. (b) Passive indoors imaging scenario, where received temperature from threat and subject presents a low contrast due to the room's warm surroundings.	11
1.7	Block diagram of a generic single-antenna FMCW architecture.	13
1.8	Time-frequency diagram of sawtooth FMCW operation.	15
1.9	(a) Beat frequency as a function of the range to the target and chirp time, when the transmitted bandwidth is 10 GHz. (b) Beat frequency as a function of the range to the target and bandwidth, when the chirp time is 1 ms.	15
1.10	(a) Range resolution as a function of the transmitted bandwidth, when $\alpha = 1$. (b) Necessary bandwidth for a range resolution of 1 cm when $\alpha > 1$	16
1.11	(a) Handheld W-band sensor, reproduced from [100]. (b) The commercial ProVizion2 portal imager from L3 Security & Detection Systems, reproduced from [101]. (c) 300 GHz standoff imager, reproduced from [102].	17
1.12	(a) Single TRX scanning where the reflector must be mechanically moved to illuminate each spot. (b) Multi-TRX array scanning where beamforming techniques are used to illuminate each spot without need for movement. (c) Combination of mechanical and array scanning, where the TRXs in the array simultaneously scan a full row of the spot grid while a mechanical movement of the reflector provides coverage for the vertical dimension.	19
1.13	(a) Vertical raster scanning. (b) Diagonal scanning. (c) Elliptical scanning.	20
2.1	Photograph of the TX and RX mm-wave chains of the 300 GHz imaging radar.	29
2.2	Spectral comparison of a focused and unfocused metallic sphere with the radar using the BEGRS reflector system.	30
2.3	Photograph detailing the reflector subsystem, diagram of the mirror's movement, and representation of the scan pattern over a human target, with the axes used throughout the text. The z axis goes through the torso, from front to back.	31

2.4	Block diagram of the 300 GHz imaging radar, divided into its main subsystems.	32
2.5	CUDA thread (<i>left</i>) and memory (<i>right</i>) organization. Reproduced from [143]	35
2.6	Main stages of the initial signal processing algorithm, in execution order from left to right, and top to bottom.	36
2.7	Scan pattern of the GMR's imaging radar and sample results from some of the signal processing events in the algorithm. Notice that the x -axis in (c) has already been labeled according to R instead of frequency, through the relationship seen in (1.18). The final product is a 3-D point cloud of the internal layer of a female mannequin wearing a T-shirt and hiding a simulated explosive device in the abdomen area. Colorbar: Range in meters (R -axis). .	37
2.8	Range gating (R -axis) in a target's silhouette.	38
2.9	Flow diagram of the peak detection algorithm for extracting range layers from a chirp's spectral samples.	39
2.10	Diagram of the elevation and rotation axes, as well as angles, in the radar's scanning mirror. Reproduced from [125].	40
2.11	Sample point cloud and reconstructed images of the internal layer of a female mannequin wearing a T-shirt and hiding a handgun in the abdomen area. Colorbar: Range in meters (R -axis).	41
2.12	Relationship between data and thread structures in the parallel algorithm for a sample image. Each point cloud spot represents one LFM ramp, and is assigned to one GPU thread. Thread blocks are formed by grouping these spots in sub-images, as seen in the sample dotted squares.	41
2.13	Sequence of events and transfers, in chronological order from left to right, from the reception of the beat frequency until the display of its image. The legend details the connections in each stage, and the color code separates hardware platforms.	42
2.14	Timeline of synchronization between scanning, image processing, and display in a sample scan. In this example, image refresh is limited by the mechanical scanning; processing can be contiguous and display can be faster, if the ADC data of all images are available sooner on the GPU.	42
2.15	Average GPU-based signal processing rate per 6000-LFM ramp image for various ADC data sizes. Configurations vary depending on signal processing and display choices seen in Table 2.4. Each configuration time is averaged over 100 images, per datapoint in the x -axis. (a) Shows the required GPU processing time per image. (b) Shows the processing speed in images per second [the inverse of (a)].	43
2.16	GPU processing time per image for all kernels, averaged over 100 images. Only one peak detection algorithm is chosen per image. FFTs are done in batched mode with CUDA's CUFFT library, while all other kernels are custom-made. Profiling is done with NVIDIA's Visual Profiler.	44
2.17	GPU resource use for all kernels, averaged over 100 images. Only one peak detection algorithm is chosen per image. FFTs are done in batched mode with CUDA's CUFFT library, while all other kernels are custom-made. Profiling is done with NVIDIA's Visual Profiler.	48
2.18	Point cloud image of a mannequin with a simulated explosive device in the abdomen area, in VTK (left) and MATLAB (right).	49
2.19	Multi-thread image refresh structure using VTK. Separate CPU threads are assigned to data transfers and image display in order to enable real-time scanning and display.	49
2.20	Histogram of recorded display times for images, measured over 100 images, with varying number of CPU threads available.	50

2.21	In black, main stages of the improved signal processing algorithm, in execution order from left to right. In red, stages removed from the previous architecture in Fig. 2.6.	52
2.22	Average GPU-based signal processing rate per 6000-LFM-ramp image for various ADC data sizes. Configurations vary depending on the signal processing and display choices seen in Table 2.7. Each configuration time is averaged over 100 images, per datapoint in the x -axis. (a) Shows the required GPU processing time per image. (b) Shows the processing speed in images per second [the inverse of (a)].	53
2.23	GPU processing time per image for the remaining kernels, averaged over 100 images. Only one peak detection algorithm is chosen per image. FFTs are done in batched mode with CUDA's CUFFT library, while all other kernels are custom-made. Profiling is done with NVIDIA's Visual Profiler.	54
2.24	GPU resource use for all kernels, averaged over 100 images. Only one peak detection algorithm is chosen per image. FFTs are done in batched mode with CUDA's CUFFT library, while all other kernels are custom-made. Profiling is done with NVIDIA's Visual Profiler.	55
2.25	Histogram of recorded display times for images, measured over 100 images, with varying number of CPU threads available.	56
2.26	Possible timeline of image processing and display, with parallelized GPU processing and host-to-device memory transfers with the use of two different CUDA streams. ADC data samples of all images must be at least synchronized with GPU transfers and at least 6 CPU threads are necessary.	58
2.27	Photograph of (a) the real-time scanning set-up for a static metallic object, and (b) the background target used (a rectangular metal frame holding a thin bar with a sphere attached at one end).	60
2.28	Sample images from a continuous video sequence at ~ 1 fps of a male entering and exiting the FoV (while extending and closing his arms), as captured by the radar. The static metal frame from Fig. 2.27(b) is present in the background. The chronological sequence starts from the top left corner, left to right, and is also indicated by the numerical order attached to each image. Colorbar: Received power in dBm.	61
2.29	Sample images from a continuous video sequence at ~ 1 fps of a male entering and exiting the FoV (while crossing his arms in front), as captured by the radar. The static metal frame from Fig. 2.27(b) is present in the background. The chronological sequence starts from the top left corner, left to right, and is also indicated by the numerical order attached to each image. Colorbar: Range in meters.	62
2.30	Sample images from a continuous video sequence at ~ 2 fps of a male entering the FoV (with his arms extended), as captured by the radar. The static metal frame from Fig. 2.27(b) is present in the background. The chronological sequence starts from the top left corner, left to right, and is also indicated by the numerical order attached to each image. Colorbar: Range in meters.	63
2.31	Illustrations of different imaging architectures according to the number of radiating elements and their placement. (a) Shows a standoff monostatic case, (b) shows walkthrough panel arrays, and (c) shows a standoff multistatic scenario.	64
2.32	(a) Specular reflection from a smooth surface, where $\theta_i = \theta_r$ (b) Diffuse reflection from a rough surface, where $\Delta\theta_r$ contains various reflection angles.	65
2.33	Standoff multi-bistatic imaging configuration, simplified from the multistatic architecture in Fig. 2.31(d).	66

2.34	Block diagram of the implemented multi-bistatic architecture using an independent TX. To generate multiple bistatic images, the TX can be moved throughout the scene.	66
2.35	Separation of the 300 GHz TX chain from the reflector system with a photography tripod for multistatic measurements.	67
2.36	(a) Bistatic FMCW radar configuration with one TRX and one TX. (b) Sawtooth FMCW signals in the bistatic configuration.	68
2.37	Diagram of the processing pipeline to obtain and combine simultaneous monostatic and bistatic point clouds. Colorbars: Range in meters.	69
2.38	(a) Measurement target consisting of a smooth metallic sheet held with a rotational positioner. (b) EPS foam slab in the same positioner. (c) Returned power from both targets in a monostatic TRX, and from TX placements at $\theta = 6^\circ$ and $\theta = 12^\circ$, as the target is being rotated around its vertical axis. (d) Zoomed-in view of the EPS results, with the metallic data removed for better visibility.	70
2.39	Monostatic images of a smooth metallic sheet (a) oriented towards the TRX, (b) rotated 3° , and (c) rotated 6° . Monostatic images of a EPS block (d) oriented towards the TRX, (e) rotated 3° , and (f) rotated 6° . Colorbar: Returned power in dBm.	72
2.40	(a) Measurement target consisting of a replica Colt 1911 handgun held with a rotational positioner. (b) Imaging result when the target is oriented towards the radar. (c) Imaging result when the target is rotated just 3° . (d) Bistatic image of the handgun when radiated with a TX with $\theta = 6^\circ$ of angular diversity. (e) Multistatic image, combined from the previous monostatic and bistatic results. Colorbar: Received power in dBm.	73
2.41	(a) Matrix of 2.5-mm-radius metallic spheres. (b) Monostatic image. (c) Bistatic image of $\theta = 3^\circ$ angular diversity. (d) Bistatic image of $\theta = 6^\circ$ angular diversity. Colorbar: Received power in dBm. The total range of the colormaps is 9 dB in all images, but not set to common limits to account for different transmit power in bistatic measurements.	74
2.42	(a) Nude female mannequin. (b) Monostatic image when the mannequin is facing the reflector system. (c) Multistatic combination of all monostatic and bistatic images when the TX is placed at four locations achieving $\theta = \pm 3^\circ$ and $\theta = \pm 6^\circ$ bistatic angles. Colorbar: Range in meters.	75
2.43	Intermediate combinations from the imaging scenario in Fig. 2.42. (a) Monostatic only. (b) Monostatic + Bistatic at $\theta = -3^\circ$. (c) Monostatic + Bistatic at $\theta = \pm 3^\circ$. (d) Monostatic + Bistatic at $\theta = \pm 3^\circ$ and at $\theta = -6^\circ$. (e) Monostatic + Bistatic at $\theta = \pm 3^\circ$ and $\theta = \pm 6^\circ$ Colorbar: Range in meters.	76
2.44	(a) Structured-light point cloud of a nude female mannequin. Colorbar: Range in meters. (b) Registered radar and structured-light point clouds of a nude female mannequin.	77
2.45	(a) Monostatic ray-tracing simulation result when a metallic sheet is oriented towards the radar. (b) Monostatic ray-tracing simulation result when a metallic sheet is rotated just 3° . Colorbar: Received power in dBm.	80
2.46	(a) 3-D model of a MP-443 handgun used as simulation target. (b) Monostatic ray-tracing simulation result when the handgun is oriented towards the radar. (c) Monostatic ray-tracing simulation result when the handgun is rotated just 3° . Colorbar: Received power in dBm.	80
2.47	Comparison between simulated and measured returned power from a smooth metallic sheet in monostatic and bistatic (at $\theta = 6^\circ$ and $\theta = 12^\circ$ angular diversities) placements, as the sheet is being rotated around its vertical axis. Measured data comes from Fig. 2.38.	81

2.48	(a) Simulated monostatic imaging result equivalent to Fig. 2.40(b). Simulated monostatic imaging result equivalent to Fig. 2.40(c) (a 3° rotation). (c) Simulated bistatic imaging result equivalent to Fig. 2.40(d) (a 3° rotation). Colorbar: Normalized received power in dB.	81
2.49	(a) Simulated monostatic image equivalent to Fig. 2.42(b). (c) Simulated multistatic combination equivalent to Fig. 2.42(c). Colorbar: Range in meters.	82
2.50	(a) Female mannequin wearing a T-shirt to hide a ceramic knife near the abdomen. (b) Placement of the knife underneath the T-shirt. (c) Monostatic image when the mannequin is facing the reflector system. (d) Monostatic image when the mannequin is rotated 10° counterclockwise. (e) Combined image of the mannequin from the bistatic result when radiated with a TX with $\theta = 20^\circ$ of angular diversity and the monostatic result in (d). Colorbar: Range in meters.	84
2.51	(a) Female mannequin concealing the gun from Fig. 2.40 with a T-shirt. (b) Photograph of the gun placement near the abdomen. (c) Monostatic image when the mannequin is facing the reflector system. (d) Monostatic image when the mannequin is rotated 10° counterclockwise. (e) Combined image of the mannequin from the bistatic result when radiated with a TX with $\theta = 20^\circ$ of angular diversity and the monostatic result in (d). Colorbar: Range in meters.	84
2.52	(a) Male human wearing a sweater to hide the handgun from Fig. 2.40(a) near the abdomen. (b) Monostatic image when the human is facing the reflector system. (c) Monostatic image when the human is rotated 10° counterclockwise. (d) Combined multistatic image of the male human from the bistatic results when radiated with TXs with $\theta = 15^\circ$, $\theta = 20^\circ$, and $\theta = 25^\circ$ angular diversities and the monostatic result in (c). Colorbar: Range in meters.	85
2.53	Histograms of the registered processing times over 1000 iterations for the 3-D Delaunay triangulation of two radar images, when performed on the GTX 970 using <i>gStar4D</i>	86
2.54	Photographs of four interconnected NVIDIA TITAN GPUs (a) and two AMD Radeon GPUs (b). The NVIDIA models are connected with an SLI bridge, while the AMD devices are connected with a CrossFire bridge.	87
2.55	FoV scan patterns for a (a) high-quality image (8000 rpm, 827 steps/s, $150 \mu\text{s}$ per chirp), (b) the fastest scan currently available (8000 rpm, 1600 steps/s, $150 \mu\text{s}$ per chirp), (c) a modified maximum (16000 rpm, 1800 steps/s, $75 \mu\text{s}$ per chirp), (d)-(f) and their corresponding imaging results for a nude female mannequin.	89
3.1	Photographs of the commercial FMCW radar kit from Silicon Radar GmbH (a) without and (b) with its dielectric lens attached. Modified from [194].	92
3.2	(a) Block diagram of the architecture of the modified 122 GHz radar from a Silicon Radar GmbH commercial kit, and (b) photograph detailing its components.	94
3.3	Photographs of the (a) TRA_120.01 and (b) TRA_120.012 MMICs from Silicon Radar GmbH, reproduced from [194].	95
3.4	SNR analysis using the MMIC-based sensor at 122 GHz for a 70 dBsm corner reflector, according to (a) (1.23) and (b) (1.24).	97
3.5	(a) Scaling of the SNR of a metallic corner reflector with a peak RCS of 70 dBsm as it is placed at increasing distances from the radar. Sweep time is $150 \mu\text{s}$, and 15.1 GHz of bandwidth are transmitted. (b) Photograph of the corner reflector used.	98

3.6	Noise floor measurements when the external ADC is connected to (a) a 50 Ω load and (b) the developed radar aiming at a millimeter-wave absorber. Varying chirp times are considered and used during the FFTs that produce the represented spectrums.	99
3.7	SNR analysis using the MMIC-based sensor at 122 GHz for a collection of common imaging objects, according to (1.23), when the dwell time is (a) 75 μ s, (b) 150 μ s, and (c) 1500 μ s.	100
3.8	(a) Sensitivity measurement and noise floor comparisons, (b) and zoom, for a metallic sphere with a 2 cm diameter illuminated by the 122 GHz radar with a 6.3 GHz bandwidth and a 150 μ s chirp. (c) Photograph of the measurement set-up.	100
3.9	Fraunhofer distance according to frequencies delimiting a 18 GHz bandwidth around the developed radar's 122 GHz center frequency.	101
3.10	Spot size estimations for the MMIC-based sensor at 122 GHz according to (3.5), with (a) and without (b) using the focusing lens.	102
3.11	Transmitted bandwidth measurement by placing a high-frequency probe in the feedback loop of the front end's PLL.	103
3.12	Maximum available bandwidth measurement for the (a) TRA_120_01 and (b) TRA_120_012 MMICs. To perform the measurement a probe is placed at the divided signal path of the PLL loop, and the frequency axis is converted to the RF range.	103
3.13	Range resolution measurement for the (a) TRA_120_01 and (b) TRA_120_012 MMICs. To perform the measurement two targets (one cardboard and one metal) are placed facing the radar, one in front of the other, and their spectrums captured. The resolution is defined as the minimum distance that still ensures a 6 dB null between targets in the spectrum.	104
3.14	Measurement of the radar's deviation from (1.18) in its beat frequency to range relationship. A 70 dBsm corner reflector is placed at increasing distances and its beat frequency is registered and compared against the theoretical value from (1.18), for the (a) TRA_120_01 and (b) TRA_120_012 MMICs.	105
3.15	Measurement of the radar's ranging accuracy. A 70 dBsm corner reflector is placed at increasing distances and its measured range according to (1.18) is compared against the actual distance to the radar, for the (a) TRA_120_01 and (b) TRA_120_012 MMICs.	105
3.16	Measurement of the radar's range-measuring accuracy when two targets in the same line of sight are present in the scene. The first target is placed at 2 m, while the second target progressively moves farther away. For the error calculation, the measured distance between targets is compared against the actual distance, for the (a) TRA_120_01 and (b) TRA_120_012 MMICs. . . .	107
3.17	Distance error for different phase fluctuations and frequencies, according to (3.6). For each frequency and fluctuation combination, target movements smaller than the corresponding curve cannot be measured accurately.	107
3.18	(a) Photograph of the positioning stage and metallic plate. (b) Distance error from phase variation, according to (3.6) for the static metallic plate during a 3 second capture using 150 μ s chirps. (c) Histogram of the distance error. . .	108
3.19	Movement sequence with the XY stage, as captured by the radar with 150 μ s chirps. In (a) the plate is moved back 10 μ m, the minimum resolution of the stage. In (b) the plate is moved back another 50 μ m. In (c) the plate is moved 1.5 mm in approximately 17 smaller steps.	109
3.20	(a) Distance error from phase variation, according to (3.6) for the static metallic plate during a 0.7 second capture using 600 μ s chirps. (b) Histogram for the distance error.	109

3.21	Distance error from phase variation, according to (3.6) for the static metallic plate during 1 s captures, taken every 5 minutes, using 600 μ s chirps.	110
3.22	IR images of the 122 GHz radar's PCB, taken every 5 minutes, as it is in continuous operation. In a span of 15 minutes the radar's maximum temperature increases by 2 °C. As can be seen, the USB connection of the microcontroller (which feeds the power consumption of the entire radar) and the MMIC are the two hottest components. Colorbar: Temperature in °C.	110
3.23	(a) Photograph of the temperature control set-up using PC ventilation fans. (b) Distance error from phase variation, according to (3.6) for the static metallic plate during 1 second captures using 600 μ s chirps, with and without temperature control. (c) Histograms of the distance errors with and without temperature control.	111
3.24	Comparison, with and without temperature control, of returned power variations from a static metallic plate using 600 μ s chirps with 6.3 GHz of bandwidth for the TRA_120.01 MMIC and 15.1 GHz of bandwidth for the TRA_120.012 MMIC.	112
3.25	Variations in the calibration procedure's performance in the 300 GHz imaging radar, when aiming at a static target, as time passes. The same reference signal is used for calibration in all times.	113
3.26	Spectral comparison of non-calibrated, calibrated, and ideal sines from a metallic sphere placed at 1 m, for the (a) TRA_120.01 MMIC transmitting 6.3 GHz of bandwidth in 150 μ s chirps and the (b) TRA_120.012 MMIC transmitting 15.1 GHz of bandwidth in 150 μ s chirps.	113
3.27	Diagram of simultaneous active, passive, and optical scanning of a human target.	114
3.28	(a) Representation of the hybrid imager's reflector system including the active, passive, and optical sensors. (b) Photograph of the side-by-side passive (gold) and active (red) sensors in front of the metallic mirror.	115
3.29	(a) Exterior and (b) internal photographs of Wavecam's passive imager, (c) and diagram of the FoV scan achieved with the combined movements of the reflector system.	116
3.30	Ray-tracing illustration showing the vertical FoV sweep performed in a single mirror turn in four steps in chronological order, from (a) to (d).	117
3.31	(a) Passive sensor's output when terminated with a 50 Ω load at 25°. (b) Histogram of the sensor's output.	118
3.32	Block diagram of the PCB used to condition and sample the passive sensor's output.	118
3.33	Outputs at different stages of the PCB diagrammed in Fig. 3.32 when a 0.9 V - 1.1 V sinusoid is input for testing and calibration.	119
3.34	Sequence of passive RX output measurements over time, as the target (a water bottle initially at 0 °C) gradually heats up to room temperature (22 °C). The X-axis marks the minimum temperature registered in the bottle.	120
3.35	Histograms of the passive RX output for four different scenarios (a mm-wave absorber, a human hand, a room-temperature water bottle, and a clear sky). The temperature of each target is listed as reference.	120
3.36	Measurement of three different targets (mm-wave absorber, human hand, and cold water bottle), where each one has 10 datapoints in the graph. The absorber was at 30 °C, the hand at 33 °C, and the bottle at 20 °C.	121
3.37	SNR of the passive RX throughout the sequence illustrated in Fig. 3.34. The SNR is computed as the logarithmic scale ratio of the mean signal in each measurement to its standard deviation.	122

3.38	Comparison of raw passive images of a male subject under various atmospheric conditions. In (a) the target is scanned in an outdoors environment, in (b) the target is scanned indoors in uncontrolled conditions, in (c) the target is scanned indoors in uncontrolled conditions in front of an open window facing the sky, and in (d) the target is scanned indoors in controlled conditions. . .	123
3.39	(a) IR camera snapshot of a male target wearing a loose T-shirt and hiding an IED threat (at 7 °C). Colorbar: Temperature in degrees. (b) Raw output image of the passive sensor. (c) Result of thresholding and filtering the raw image with a 2-D asymmetric median kernel. (d) Final passive image, after reducing the image value interval. (e) Graph of a single column in (c), cutting through background, body, and IED. The samples corresponding to the IED can be clearly separated from the body.	123
3.40	(a) IR camera snapshot of a male target hiding a threat (a filled room-temperature 1 L water bottle) underneath a thick sweatshirt and T-shirt. Colorbar: Temperature in degrees. (b) Filtered and thresholded output image of the passive sensor.	124
3.41	Measurement sequence with the passive sensor of a male subject with a rectangular target in his chest, as the room cools down from air conditioning. The sequence shows both 94 GHz and IR images. Colorbar: Temperature in degrees. The image in (a) is taken at the beginning of the sequence, with no air conditioning. In (b) the air conditioning is turned on, and images approximately every 30 minutes are taken in (c) - (f).	125
3.42	Measurement sequence with the passive sensor as a threat warms up, while hidden beneath a male subject's T-shirt. The threat is the previously-used IED, made up of three individual capacitors. The left column shows IR close-ups of the torso, while the center and right column are the raw and processed images obtained with the passive RX, respectively. Colorbar: Temperature in degrees.	126
3.43	Histograms of the raw RX outputs when the sensor aims at a water bottle in various concealment scenarios, and photographs of some of the set-ups. Outputs from the unconcealed bottle and a mm-wave absorber are also included for comparison. In (a) the concealed bottle is covered with a polyester T-shirt, in (c) with a cotton sweater, and in (e) with denim jeans. The photograph in (b) shows the uncovered bottle, the one in (d) shows the mm-wave absorber, and the one in (f) shows the covered bottle (with denim).	127
3.44	Passive images of a cool water bottle in the concealment scenarios analyzed in Fig. 3.43. In (a) the bottle is uncovered, in (b) the bottle is covered with a polyester T-shirt, in (c) the bottle is covered with a cotton sweater, and in (d) the bottle is covered with denim.	128
3.45	Passive image data used for added noise testing. The subject is a male wearing a T-shirt hiding a cool water bottle in the abdomen. (a) Raw samples of the passive sensor's output in digital values, (b) histogram of raw and filtered samples, (c) raw image, and (c) filtered image.	129
3.46	Passive image data from Fig. 3.45 with additional AWGN. The left column are histograms of the additional AWGN introduced in each row, the center column is the corresponding raw image, while the right column is the filtered image. The top row is a scenario where the AWGN has a standard deviation equal to 50% of the standard deviation of the data in Fig. 3.45(a), the middle row is a scenario where the AWGN is raised to 100% of the original standard deviation, and the bottom row is a scenario where the AWGN is 150% of the original standard deviation.	130

3.47	(a) Measurement of the cross-range resolution of the passive sensor along the horizontal dimension using the reflector system. 2 cm-diameter frozen plastic spheres are placed 5 cm apart, and sequentially brought together in front of a human. The distances above the images represent the separation between the objects' centers. (b) Same sequence as above, for the vertical dimension. . . .	131
3.48	Measurement of the spot size of the passive sensor along the horizontal (X) and vertical (Y) dimensions using the reflector system. A 2 cm-diameter frozen plastic sphere is illuminated, and single line scans in each dimension are measured.	132
3.49	(a) Photograph of the male subject (the location of the threat underneath the T-shirt is indicated). (b) Scans of the set-up on the left, using the modified active sensor in the reflector system. In the examples on the top the person is hiding an IED composed of three cylindrical capacitors, while on the bottom figures no threat is hidden. Examples on the left are representations in 3-D point cloud format while those on the right depict a surface reconstruction of the point clouds. A photograph of the hidden threat is inset between the examples. Colorbar: Range in meters.	133
3.50	Noise power level as a function of the maximum reflected power per pixel for all pixels obtained from Fig. 3.49(b) (with the IED). The red data uses a single range bin near the target to estimate noise power, while the blue data uses the mean power of 50 range bins around the target.	133
3.51	Measurement of the cross-range resolution of the active sensor using the reflector system. The top row is the resolution analysis in the horizontal dimension, and the bottom row in the vertical dimension. The centers of two metallic spheres are placed at a varying distance [(a) and (d) 4 cm, (b) and (e) 3 cm, (c) and (f) 2 cm]. Colorbar: Received power in dBm.	134
3.52	Measurement of the cross-range resolution of the active sensor along the horizontal and vertical dimensions using the reflector system. 5 mm-diameter metallic spheres are placed 7.5 cm apart, in both directions, over a polystyrene sheet. Colorbar: Received power in dBm.	135
3.53	Measurement of the spot size of the active sensor along the horizontal (X) and vertical (Y) dimensions using the reflector system. A 5 mm-diameter metallic sphere is radiated, and single line scans in each dimension are measured. . . .	136
3.54	Interference in the passive sensor's output caused by the active sensor. (a) Measurement set-up, and passive sensor output comparisons in the (b) time and (c) frequency domains.	137
3.55	Passive images of a male subject with no hidden threat. (a) Active sensor turned off, (b) active sensor transmitting chirp (interference boxed in red), and (c) comparison of the passive sensor's raw output in both cases.	138
3.56	Dependence of the interference in the passive sensor with the active sensor's transmitted bandwidth.	139
3.57	Imaging example when the active sensor is off-center in the reflector system. (a) Passive image with interference shifted slightly away from the subject's abdomen (interference boxed in red), and (b) corresponding active point cloud. Colorbar: Range in meters.	139
3.58	Passive and active images of a male subject with no hidden threat when the active sensor transmits pulsed chirps in a 1-to-10 ratio. (a) Raw passive image, (b) raw active image, and (c) interpolated and filtered active image. Colorbar: Returned power in dBm.	140
3.59	Passive sensor outputs under different active sensor transmissions. (a) Output in the time domain, and (b) histograms of those outputs.	141
3.60	Sample combination of active sensor (a) and Kinect (b) outputs via an overlay (c) of their two images with 50% transparency.	141

3.61	Diagram of the image combination process of all three sensors involved in the imager.	142
3.62	(a) Photograph from the Kinect's RGB camera, with face blurring applied and the detected skeleton represented as red segments. (b) Output of the Kinect's IR sensor with a grayscale colormap representing depth. (c) Initial depth point cloud with RGB texture from the camera. (d) Final Kinect 3-D point cloud, later used as depth mask, after segmenting the initial point cloud within the subject's range window.	142
3.63	Breakdown of the processing steps for fusing two images using the DWT-based technique proposed. The unfused input images are from the Kinect and active sensor, for the imaging scenario in Fig. 3.49. Decompositions that include passing through the high-pass filter are reduced to simply the borders of the input image, because these boundaries produce a sharp change in image data (interpreted as a high-frequency component). This is also why the top decomposition, created by passing only through the low-pass filter, is approximately the same image as the input; values inside the image's borders have very smooth changes (interpreted as low-frequency information). The Kinect's depth data has been used to generate image $I_1(x, y)$, but the colored range information in the final fused image $[Z(x, y)]$ comes from the active sensor. .	145
3.64	Diagram of the complete processing chain and outputs at different stages. Blocks in different colors are processed in separate CPU threads.	146
3.65	(a) Imaging scenario of male target hiding a room-temperature IED of three cylindrical capacitors underneath a T-shirt in the abdomen area. In this case both sensors expose the threat. (b) Range-based representation of the active sensor data. (c) Active sensor data fused with the Kinect point cloud. Colorbar: Range in meters. (d) Passive sensor data. (e) Passive sensor data fused with the Kinect point cloud.	147
3.66	(a) Range point cloud of the active sensor. Colormap: Range in meters. (b) Registration of the range point cloud with the Kinect's point cloud. (c) Filtered grid with range data. (d) Fused range data and Kinect mask. (e) Returned power point cloud of the active sensor. Colormap: Returned power in dBm. (f) Registration of the returned power point cloud with the Kinect's point cloud. (g) Filtered grid with returned power data. (h) Fused returned power data and Kinect mask.	147
3.67	(a) Thresholded and filtered output of the passive sensor. (b) Cropped passive image, where the red region corresponds to the result of performing a colored flood-fill of holes inside the torso. (c) Fused passive data and Kinect mask. .	148
3.68	(a) Imaging scenario of male target hiding a room-temperature pocketknife in the left pocket of his jeans. (b) Range-based representation of the active sensor data. (c) Active sensor data fused with the Kinect point cloud. Colorbar: Range in meters. (d) Passive sensor data. (e) Passive sensor data fused with the Kinect point cloud. This is the only representation in this scenario where the threat is exposed.	149
3.69	(a) Imaging scenario of male target hiding a simulated plastic explosive with shrapnel, with has warmed up to body-temperature, in the right side of his abdomen. (b) Range-based representation of the active sensor data. (c) Active sensor data fused with the Kinect point cloud. Colorbar: Range in meters. This is the only representation in this scenario where the threat is exposed. (d) Passive sensor data. (e) Passive sensor data fused with the Kinect point cloud.	149

3.70	(a) Imaging scenario of male target hiding a plastic 1 L sport bottle full of room-temperature water underneath a T-shirt in his stomach area. In this case both, sensors expose the threat. (b) Range-based representation of the active sensor data. (c) Active sensor data fused with the Kinect point cloud. Colorbar: Range in meters. (d) Passive sensor data. (e) Passive sensor data fused with the Kinect point cloud.	150
3.71	(a) Diagram of array scanning with two sensors, placed on a moving platform, simultaneously transmitting at different parts of the FoV. (b) Photograph of the four passive RXs mounted in an array configuration in the reflector system.	150
3.72	Multi-sensor scanning of a male subject with four passive RXs in an array configuration. (a) Joined raw outputs of the 4 sensors, (b) filtered outputs, and (c) normalized and thresholded outputs.	151
3.73	Histograms of the raw outputs of the four passive RXs during the imaging scenario in Fig. 3.72.	152
3.74	Photographs of two 122 GHz radars placed in an array configuration for simultaneous scanning in the imager.	152
3.75	Multi-sensor scanning of a male subject with two 122 GHz radars in an array configuration. (a) Point cloud from the first sensor, (b) point cloud from the second sensor, (c) and combined point clouds. Colorbar: Range in meters. (d) Point cloud from the first sensor, (e) point cloud from the second sensor, and (f) combined point cloud. Colorbar: Returner power in dBm.	153
3.76	Combined multi-sensor scans of a male subject with two 122 GHz radars in an array configuration, after processing each sensor's images with different detection thresholds. (a) Colorbar: Range in meters. (b) Colorbar: Returned power in dBm.	153
3.77	(a) Set-up for hybrid imaging with the 94 GHz RX mounted in the 300 GHz reflector system. (b) Interferences in the 94 GHz RX caused by the 300 GHz radar's TRX.	154
3.78	(a) Nude female mannequin used as target, (b) active image of the mannequin with the previous PCB, and (c) active image for that same scenario with the new PCB. Colorbar: Returned power in dBm.	155
4.1	Sample display from a mm-wave security scanner from L-3 Technologies running ATD algorithms.	158
4.2	Sample images from the DHS dataset. (a) Front view of a man with no threat. (b) Side view of a woman with no threat. (c) Side view of a man with a concealed threat near his right knee. (d) Side view of a woman with a concealed threat near her left calf. (e) Side view of a man with a concealed threat near his right ribcage. (f) Back view of man with a concealed threat near his shoulder blades.	159
4.3	Sample 360° scan of a male subject in the DHS dataset.	160
4.4	Diagram of possible threat locations for the images in the DHS dataset. Reproduced from [238].	161
4.5	Flow diagram of a perceptron.	163
4.6	Generic diagram of a simple CNN with few layers.	163
4.7	Commented diagram of AlexNet's architecture, based on an illustration in [263]. Notice that the network's flow has two stacked layers dividing processing of the input image in two halves, due to the original authors' use of two GPU's during development.	164
4.8	Examples from a single scenario of the DHS dataset where certain body areas are not visible in some viewpoints. (a) Shows viewpoint number 9, where body area 5 (the chest) is not visible, and (b) is the complimentary case; viewpoint 1 where body area 17 (the upper back) is not seen.	168

4.9	Probability of threat existence in each body zone for the labeled images in the DHS dataset.	168
4.10	Sample comparison between an original (a) and processed (b) image from one of the DHS dataset's scenarios.	169
4.11	Average training duration in minutes for each network in the model.	170
4.12	Evolution of training accuracy for three of the model's CNNs.	171
4.13	Diagram of a generic confusion matrix for a 2-class classifier.	172
4.14	Confusion matrices for all 17 CNNs in the developed model, in sequential order, left to right and top to bottom. The vertical axis in each chart represents the predicted class, while the horizontal axis represents the actual class. Matrix layout follows the diagram in Fig. 4.13.	173
4.15	Evolution of training loss for three of the model's CNNs.	174
4.16	Training loss for each CNN in the developed model.	174
4.17	Weight representation for the first convolutional layer of all 17 CNNs in the developed model, in sequential order, left to right and top to bottom.	176
4.18	Sample processed image from the DHS dataset showing a person with an oval threat in the upper chest area (body region 5). Threat boxed in blue, only for this display	177
4.19	Activation maps for network 5's first convolutional layer when input with the image in Fig. 4.18.	178
4.20	Comparison of the image from Fig. 4.18 and the maximum activation map from Fig. 4.19.	179
4.21	Activation maps for network 5's fifth convolutional layer when input with the image in Fig. 4.18.	180
4.22	Comparison of the image from Fig. 4.18 and the maximum activation map from Fig. 4.21.	180
4.23	(a) Comparison of the image from Fig. 4.18 and the maximum activation map from the network's ReLu layer after the fifth convolutional layer, (b) along with an overlay of the two after coloring the activation map in red.	181
4.24	Comparison of a processed image from the DHS dataset showing a person with no threat in the upper chest area (body region 5) and the maximum activation map from the network's ReLu layer after the fifth convolutional layer.	181
4.25	Threats included in the IR dataset, in order from left to right: ice pick, 9 mm handgun, ceramic knife, shrapnel explosive (plasticine+bolts), water bottle, IED (capacitors).	183
4.26	Sample frames from an imaging scenario in the IR dataset where a male subject hides a cool water bottle underneath a T-shirt. The frames show the four $\sim 90^\circ$ viewpoints the subject is asked to maintain during one minute each. Colorbar: Temperature in degrees.	184
4.27	Sample frames from the scenario in Fig. 4.26 taken just 1 second apart. Frame (a) is labeled as a threat scenario because it displays the full bottle, while frame (b) is labeled a no threat scenario because it only shows part of the object. Colorbar: Temperature in degrees.	184
4.28	(a) Raw IR frame and (b) cropped result. In the cropped example all unnecessary shapes in the image have been erased. Colorbar: Temperature in degrees.	185
4.29	Distribution of images from the IR dataset into the 7 threat classes considered.	185
4.30	Average image of the IR dataset generated for image diversity calculation.	186
4.31	Training (a) accuracy and (b) loss for the CNN when inputting images of a single subject hiding cooled threats placed in his torso underneath a T-shirt.	187
4.32	Diagram of a generic confusion matrix for a N -class classifier focusing on class c_k	187

4.33	Confusion matrices for (a) training and (b) validation for the CNN when inputing images of a single subject hiding cooled threats placed in his torso underneath a T-shirt. The vertical axis in each chart represents the predicted class, while the horizontal axis represents the actual class. Matrix layout follows the diagram in Fig. 4.32.	188
4.34	Comparison between the two male subjects considered in the dataset for the same threat scenario (a concealed handgun).	189
4.35	Confusion matrices during testing for the CNN when inputing images of two male subjects when the training images come from (a) a single subject or (b) both subjects. The vertical axis in each chart represents the predicted class, while the horizontal axis represents the actual class. Matrix layout follows the diagram in Fig. 4.32.	190
4.36	Sample images of additional scenario variations considered in the full IR dataset, such as (a) uncooled threats, (b) thicker clothing, or (c) random threat placements. In all images the subject is hiding the same handgun. . .	191
4.37	Confusion matrices during testing for the CNN when inputing images from the full IR dataset, after it has been trained with (a) a subset or (b) the full dataset. The vertical axis in each chart represents the predicted class, while the horizontal axis represents the actual class. Matrix layout follows the diagram in Fig. 4.32.	192
5.1	Diagrams of (a) the sensor as it moves through a road with a guide installed, and (b), (c) of the plastic guide with a codification based on two levels of cavity depth. The sensor's movement (vector V) is used to scan the guide. The set-up's axes, used throughout the text, are also depicted.	199
5.2	Maximum vehicle speed for the proposed technique as a function of bit size (l_{bit}), chirp time (T_{chirp}), and number of chirps integrated per bit (n).	201
5.3	Block diagram and photograph of the RF architecture of the radar used during measurements.	203
5.4	Diagram (a) and photograph (b) of the focusing scheme for the proposed set-up. Dielectric lenses are used to focus the transmission onto an area small enough to cover only a single bit. F_1 and F_2 are the lenses' focal lengths (5 cm and 20 cm, respectively).	203
5.5	Diagram (a) and photograph (b) of the radiation scheme for the proposed set-up. The MMIC-based radar's HDPE lens is used to collimate radiation on the OUT. The distance to the OUT is set to 32.5 cm, which is the working range taking into account the focal lengths of the lenses in Fig. 5.4, to offer comparable results.	204
5.6	(a) Comparison of uncalibrated, calibrated and ideal spectrums received from a metallic sheet, as captured by the 300 GHz radar. (b) Set-up used for the measurement of the calibration's reference signal from a metal sphere.	205
5.7	Measurements at (a) 300 GHz and (b) 122 GHz (Configuration 3) of a metallic sheet, a solid slab of HDPE, and a mm-wave absorber. RX bandwidth is 20 MHz, sweep period is 150 μ s, and a Hamming window is applied.	205
5.8	XY spot size measurement at (a) 300 GHz and (b) 122 GHz (Configuration 3), with a return loss criteria, and using metallic coins of various sizes as targets.	206
5.9	Range measurement sequence at (a) 300 GHz and (b) 122 GHz (Configuration 3) with a 1€ coin moving 10 cm back and front starting from 32.5 cm, in 1 cm steps, to analyze tolerance to variations of target placement in the Z axis.	207

5.10	Measurement of losses due to the sensor's orientation shifts in the Z axis at (a) 300 GHz and (b) 122 GHz (Configuration 3). A 0° shift means the target's surface normal aims directly at the radar. For the 122 GHz radar, the measurements performed include data with and without its dielectric lens. . .	208
5.11	Spectrum of solid plastic slab in different humidity scenarios at (a) 300 GHz and (b) 122 GHz (Configuration 3). In the <i>wet</i> scenarios water is applied with a diffuser on the surface. In the <i>film</i> scenario plastic film is placed over the wet surface.	208
5.12	Spectrum of a block of road asphalt at (a) 300 GHz and (b) 122 GHz (Configuration 3), compared against the received reflections from a metallic sheet, a solid HDPE slab, and a mm-wave absorber.	209
5.13	Photograph of the first two $15 \times 2.5 \times 5$ cm ³ HDPE guides designed, with (a) internal-gap and (b) surface information.	210
5.14	Spectrum of the bit columns in Fig. 5.13(a) at 300 GHz. Columns are numbered according to their position in Fig. 5.13(a), starting at the left side of the photograph. The numbers in parenthesis are the error, produced by Z -axis resolution and FFT precision, in the measurement of distance for each surface.	211
5.15	Spectrum of the first column (starting from the left) in Fig. 5.13(a), for Configuration 3 of the 122 GHz radar.	211
5.16	Continuous scan of a solid HDPE slab at (a) 300 GHz and (b) 122 GHz (Configuration 3), with a -50 dBm threshold. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.	212
5.17	Continuous scan of the guide in Fig. 5.13(a) at (a) 300 GHz and (b) 122 GHz (Configuration 3), starting at the leftmost column, with a -50 dBm threshold. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.	212
5.18	Continuous measurement sequence for the guide in Fig. 5.13(b) at (a) 300 GHz and (b) 122 GHz (Configuration 3), starting at the rightmost column and with the bits at the back surface, with a -50 dBm threshold. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.	213
5.19	Photograph of an HDPE guide based on triangular gaps as bits and diagram of the radar orientations considered when measuring the guide.	214
5.20	Continuous measurement sequence for the guide in Fig. 5.19 at 300 GHz, with a -50 dBm threshold. In (a) the radar is placed in the first orientation, according to Fig. 5.19(b), while in (b) the radar is in the second orientation of that figure. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.	214
5.21	Guide in Fig. 5.19 for the (a) first and (b) second orientations considered, for Configuration 3 of the 122 GHz radar. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.	215
5.22	Hybrid item-labeling system in container storage, where different tagging technologies are used at different storage levels to optimize each technique's advantages.	216
5.23	(a) Photograph and (b) diagram of the measurement set-up for vital sign detection and tracking with the 122 GHz radar.	220
5.24	Block diagram displaying the signal processing stages for vital sign monitoring with the 122 GHz radar.	220
5.25	(a) Measured chest displacement during the subject's low breathing rate scenario. (c) Measured displacement during the subject's high breathing rate scenario.	221

5.26	(a) Measured chest displacement during the subject's breathing. (b) Spectral representation of the measured data.	221
5.27	(a) Measured results for a resting subject's heart rate, compared against the ECG board's data. (b) Spectral representation of the radar-measured data. .	222
5.28	(a) Diagram of the pulse's waveform as measured with a PPG, reproduced from [394]. (b) Isolated heart cycle measured with the radar in Fig. 5.27(a). .	223
5.29	(b) Measured heartbeat displacement with the 300 GHz radar. (b) Spectral representation of the measured data.	224
5.30	(b) Measured respiration displacement with the 300 GHz radar. (b) Spectral representation of the measured data.	224
5.31	Block diagram of the proposed bit-reading algorithm for the data acquisition technique.	226
5.32	Examples of information-density improvements for the data-acquisition technique. (a) shows the 2-bit-level single guide scheme, (b) shows a guide with increased bits in each X step, while (c) shows two guides stacked in the Z dimension.	226
5.33	Joint measurement of a resting male subject's breathing rate using the 122 GHz radar and a Microsoft Kinect device.	228
5.34	(a) Distance fluctuation from phase variations with the 122 GHz aiming at a metallic door that suddenly closes, and (b) spectral representation of the captured data.	230
5.35	(a) Distance fluctuation from phase variations with the 122 GHz aiming at the floor, near to where a human subject abruptly falls to the floor, and (b) spectral representation of the captured data.	230
5.36	Distance fluctuation from phase variations with the 122 GHz aiming at the underside of a table as several objects are placed and removed directly on top. .	231
5.37	Spectral representation of the distance fluctuation from phase variations with the 122 GHz aiming at a speaker as sinusoidal signals at (a) 10 Hz, (b) 100 Hz, and (c) 2000 Hz are played.	232
5.38	Block diagrams comparing multi-node architectures in (a) unsupervised and (b) supervised settings.	233
A.1	Double threshold power boundary applied to a sample chirp spectrum.	267
A.2	Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a female mannequin is hiding a handgun underneath a T-shirt, when the noise and dynamic range thresholds are (a) -89 dBm and 12 dB, (b) -79 dBm and 12 dB, and (c) -79 dBm and 30 dB, respectively. Colorbar: Range in meters.	268
A.3	Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a female mannequin is hiding a handgun underneath a T-shirt when the colormap is associated to (a) range (in m), (b) returned power (in dBm), and (c) phase (in degrees).	269
A.4	Imaging scenario for the orientation sequence measured in Fig. A.5, where a smooth metallic plate is placed 8 m away from the radar using a tripod. . . .	269
A.5	Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a smooth metallic plate is imaged with (a) no rotation, (b) 1° rotation, and (c) 2° rotation around the y -axis. Colorbar: Received power in dBm. . .	270
A.6	Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a female mannequin is hiding a handgun underneath a T-shirt when (a) one peak, (b) two peaks, and (c) three peaks are selected. Colorbar: Range in meters.	270

A.7	Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a female mannequin is hiding a handgun underneath a T-shirt when the (a) first range peak, (b) last range peak, and (c) maximum power peak are chosen. Colorbar: Range in meters.	271
A.8	Collection of representations available in the 300 GHz imaging radar for the results of scanning a female mannequin hiding a handgun underneath a T-shirt. (a) is a point cloud display of the 3-D coordinates, (b) is the result of applying a cubic interpolation of the point cloud into a rectangular grid occupying the FoV with a 1 mm step size, and (c) is the output from 2-D filtering the previous interpolated grid with a median kernel to smooth its surface. Colorbar: Range in meters.	272
A.9	Collection of representations available in the 300 GHz imaging radar for the results of scanning a female mannequin hiding a handgun underneath a T-shirt. (a) is a point cloud display of the 3-D coordinates, (b) is the result of applying a Delaunay triangulation in 3-D to the point cloud, where each triangle is colored according to its range, and (c) is that same reconstruction but also providing texture to the facets of the triangles generated. Colorbar: Range in meters.	272
A.10	Examples of interpolated and 2-D filtered outputs for a scan at 300 GHz of a female mannequin hiding a handgun underneath a T-shirt, where the filter's kernel sizes are (a) 1x1 cm ² , (b) 2.5x2.5 cm ² , and (c) 4x4 cm ² . Colorbar: Range in meters.	273
A.11	Examples of reconstructed surfaces for a scan at 300 GHz of a female mannequin hiding a handgun underneath a T-shirt, using a 3-D Delaunay triangulation, where the previous to the triangulation the point cloud's axes are compressed by factors of (a) 1.3 in x only, (b) 1.3 in x and 1.6 in y , and (c) 1.6 in y only. Colorbar: Range in meters.	274
A.12	Resulting 3-D point clouds from the 300 GHz imaging radar, after being processed with the DBSCAN clustering technique, of scenarios where (a) a male human wearing a shirt hides a plastic bottle (colored orange), (b) a female mannequin wearing a T-shirt hides a rectangular simulated explosive device (colored orange), and (c) that same mannequin hiding a gun (colored brown). The colors represent the different clusters detected.	275
A.13	Comparison of the clustering results of (a) serial and (b) parallel versions of the DBSCAN algorithm for an imaging scenario of the 300 GHz radar where a female mannequin conceals a gun (colored brown) under a T-shirt.	276
A.14	Block diagrams of the processing stages for three different ATD strategies for the 300 GHz imaging radar.	276
A.15	Outputs of the processing stages for the first threat detection strategy in Fig. A.14. In (a), corresponding to step A in Fig. A.14, the point cloud captured with the 300 GHz imaging radar for a scenario of a female mannequin wearing a T-shirt and hiding a gun is shown. In (b), corresponding to step B in Fig. A.14, the captured point cloud is registered to a reference point cloud of the mannequin with no clothing or threat. In (c), corresponding to step C in Fig. A.14, points belonging to both clouds or not protruding from the body are eliminated, leaving only those from the gun. In (d), corresponding to step D in Fig. A.14, the boundary of the extracted coordinates is computed, showing visual similarity to a gun. Colorbar: Range in meters.	277

A.16	Outputs of the processing stages for the second threat detection strategy in Fig. A.14. In (a), corresponding to step A in Fig. A.14, the point cloud captured with the 300 GHz imaging radar for a scenario of a female mannequin wearing a T-shirt and hiding a gun is shown. In (b), corresponding to step E in Fig. A.14, the captured point cloud is clustered with the DBSCAN technique. In (c), corresponding to step F in Fig. A.14, the boundary of the gun cluster is computed and visually compared against a reference threat shape. In (d), corresponding to step E in Fig. A.14, the Fourier descriptors of the boundaries are compared in a spectral representation. Colorbar: Range in meters.	278
A.17	Outputs of the processing stages for the third threat detection strategy in Fig. A.14. In (a), corresponding to step H in Fig. A.14, the interpolated and filtered grid from the 300 GHz imaging radar for a scenario of a female mannequin wearing a T-shirt and hiding a gun is shown. In (b), corresponding to step I in Fig. A.14, the grid is processed with a watershed algorithm to segment surfaces in different ranges (the segmented surfaces are highlighted with different colors). In (c), corresponding to step J in Fig. A.14, the gun segment is extracted. Colorbar: Range in meters.	279
A.18	Sample point cloud images from the initial database for the 300 GHz imaging radar. The top row is of a female mannequin (a) wearing no clothing and hiding no threat, (b) wearing a T-shirt and hiding a gun, and (c) wearing a T-shirt and hiding a simulated explosive device. The bottom row is of a male human, covering the same clothing and threat scenarios, in the same order. .	280
A.19	Block diagram comparison between (a) the original mm-wave architecture (with items to be removed in red) and (b) the compact configuration.	282
A.20	Received power for a metallic target as its range grows progressively, for the bistatic and compact architectures.	282
A.21	(a) Female mannequin with a T-shirt hiding a set of fake cylindrical explosives in the stomach area. (a) Inner layer image with the compact architecture. (b) Inner layer image with the original architecture. Colorbar: Range in meters..	283
A.22	Imaging comparison of the maximum-power layer of a nude female mannequin. (a) is the result of using the original x6 multiplier, and (b) is the result of using the updated version. Colorbar: Received power in dBm.	284
A.23	Photographs of the (a) previous and (b) new power supply units for the imaging radar.	284
B.1	Active x6 multipliers' output power for the imaging radar's 95.4 GHz - 104.4 GHz bandwidth.	286
B.2	Comparison of output power at the x6 multipliers' output with varying input power (P_{in}). (a) is the original mutliplier, and (b) is the updated version. . .	286
B.3	Isolator's output power for the imaging radar's 95.4 GHz - 104.4 GHz bandwidth.	287
B.4	MPA's output power for the imaging radar's 95.4 GHz - 104.4 GHz bandwidth.	287
B.5	Comparison of power at the MPA's output with varying input power (P_{in}) for the active x6 multipliers. (a) is the original mutliplier, and (b) is the updated version.	288
B.6	MPA's output power for the imaging radar's 95.4 GHz - 104.4 GHz bandwidth.	288
B.7	Comparison of power at the x3's output with varying input power (P_{in}) for the active x6 multipliers. (a) is the original mutliplier, and (b) is the updated version.	289
B.8	(a) Measured S_{21} parameter from Mini-Circuit's bandpass filter. (b) Band-pass zoom of the S_{21} measurement. (c) Superior stopband zoom of the S_{21} measurement.	292

C.1	(a) CAD diagram and (b) photograph of the scanning mirror, its mechanic structure, and the devices (motors and reduction drive) involved in its movement.	295
C.2	CAD render of the PCB in charge of motor control and encoder handling. . .	296
C.3	Block diagram presenting the processes and subsystems involved in motor control and positioning.	296
C.4	Diagram illustrating the devices and connections involved in the radar's real-time positioning subsystem.	297
C.5	Photographs of (a) the inductive sensors and (b) the elevation limits in the mirror's mechanical structure.	297
C.6	Chronological sequence of mirror positions during the vertical alignment. In the first step, the mirror begins the initialization in a random position. The motor is then ordered to move upwards until it reaches the upper structural limit, in the second step. In the final step, the motor has been ordered to move downwards a fixed number of steps to start the scanning in a mechanically safe position.	298
C.7	Examples of the elevation motor's output signal captured by the ADC, when the motor moves (a) at ~ 0.6 fps, and (b) at ~ 1.2 fps. When the pulse train is in its <i>high</i> state, the motor is in movement. Otherwise, the motor has stopped. Circles in red are the spots where the C++ software has detected movement starts or stops.	299
C.8	Block diagram showing the encoder output paths in the microcontroller PCB, which receives and transfers their information to the PC.	299
C.9	Histograms corresponding to 100 positioning data packets sent over (a) TCP and (b) UDP through the Ethernet connection between the PCB and the PC.	300
C.10	Comparison of the interpolated offline-positioning (in red) and the real-time 1-packet-per-ramp positioning (in blue) techniques. In the zoomed view, some of the limitations of the real-time positioning are better seen. Individual spots are lost due to microcontroller clogging, and entire sections are empty due to motor speed misestimation.	301
C.11	Chronological sequence of events in a typical continuous scan. When the mirror reaches the upper or lower scanning limits, a data packet is sent to the PC containing the initial position information for the next scan. While the elevation motor is moving up or down, no data packets are sent (positioning is done via interpolation from the initial position)	301
C.12	Comparison of the interpolated offline-positioning (in red) and the real-time 1-packet-per- n -ramps positioning (in blue) techniques. The n value is 10 in (a), 100 in (b), and 1000 in (c).	302
C.13	Registered times for two different elevation events (movement and stops) during a 100-image continuous scan, for varying elevation motor speeds. The data around 0.1 s (equal for all speeds) correspond to the duration of stops in the upper or lower limit when an individual scan is finished, and it is limited by the motor controller and its code. Little variation in time required can be appreciated during movement at any speed.	303
C.14	Histograms of registered initial rotation angles for 100 consecutive images during a continuous scan, for varying motor speed.	303
C.15	Registered initial rotation angles in chronological order for 33 consecutive images during a continuous scan, for varying motor speed.	304
C.16	Registered initial elevation angles in chronological order for 33 consecutive images during a continuous scan, for varying motor speed.	304

C.17	Examples of ADC buffer overflow during the continuous sampling of a sinusoidal signal. The red dashed rectangles indicate where samples have been lost. In (a) the sampling frequency is 40 MS/s, and only 1 ADC channel is being captured. In (b) the sampling frequency is 25 MS/s, and 2 ADC channels are being captured.	305
C.18	Data transfer for various interfaces compared against varying sampling rates, as total data size increases and 2 ADC channels are captured.	306
C.19	Example of uninterrupted capture with the ADC during continuous sampling of a sinusoidal signal. The sampling frequency is 5 MS/s and 2 ADC channels are being captured.	307
C.20	Diagram illustrating the bandpass sampling and filtering process over the received signal's spectrum, used to enable continuous scanning with the current ADC.	308
D.1	Photograph of Silicon Radar's 122 GHz FMCW radar kit.	309
D.2	Initial (as given by Silicon Radar's kit) and modified radar architectures. Items in red have been removed, items in green have been added, and items in black remain untouched.	310
D.3	Voltage ramps input into the MMIC's VCO for FMCW operation in the modified (in blue) and original sensor (in red), showing both types of off-time preventing continuous transmission in the unmodified case. In the larger graph the blue data shows that transmission is now continuous, while the zoomed-in view confirms that it is composed of voltage ramps of the same amplitude as before the modifications and with faster sweeps.	311
D.4	(a) Photograph of the top side of the new control board. (b) Photograph of the bottom side of Silicon Radar's modified RF board. (c) Photograph of the top side of Silicon Radar's modified RF board.	312
D.5	(a) Comparison of the clock signal at different places in the boards. (b) Zoom of the final clock signal used for sampling in the new ADC.	312
D.6	(a) Spectrum and (b) phase noise performance of the 75 MHz reference oscillator.	313
D.7	(a) Spectrum and (b) phase noise performance of the frequency divider. . . .	313
D.8	(a) Comparison of the end-of-ramp signal from the RF board's synthesizer, before and after passing through a comparator. (b) Zoom of the pulse associated to the downwards frequency sweep, before and after passing through a comparator.	314
D.9	Comparison of the noise floor and sampling frequency for the previous and new configuration of the active imaging sensor. A target is placed at 2 m, sweep and FFT time are 600 μ s, and 6.3 GHz of bandwidth are transmitted .	315
D.10	Evolution of MMIC temperature during continuous transmission. Temperature becomes stable after approximately 10 minutes of uninterrupted operation.	315
D.11	Photograph comparing the modified 122 GHz radar board developed during this thesis (left) and its updated version (right).	316
F.1	Photographs of the blackValue [®] sorting system, developed by the Fraunhofer Institute for High Frequency Physics and Radar Techniques. Image reproduced from [462].	326
F.2	Photographs of the PLA pieces considered for evaluation. (a) Black-colored PLA of 50% (top, <i>PLA1</i>) and 10% (bottom, <i>PLA2</i>) rectilinear fill pattern, (b) Red-colored PLA of 50% (top, <i>PLA1</i>) and 10% (bottom, <i>PLA2</i>) rectilinear fill pattern.	326
F.3	Spectrum comparison at (a) 300 GHz and (b) 122 GHz (Configuration 3) of a reference metallic sheet and the two HDPE pieces considered.	327

F.4	Spectrum comparison at (a) 300 GHz and (b) 122 GHz (Configuration 3) of a reference metallic sheet and the two black PLA pieces considered, when illuminating them on a solid side or an open side (exposing the piece's infill pattern).	328
F.5	Spectrum comparison at (a) 300 GHz and (b) 122 GHz (Configuration 3) of a reference metallic sheet and the two red PLA pieces considered, when illuminating them on a solid side or an open side (exposing the piece's infill pattern).	329
F.6	Spectrum comparison at 300 GHz of the black and red PLA designs when illuminating (a) their solid face or (b) their open face.	330
F.7	Spectrum comparison of the two PLA pieces considered, when illuminating them on a (a) solid side or an (b) open side (exposing the piece's infill pattern), for Configuration 3 of the 122 GHz radar.	330
F.8	Spectrum comparison at 300 GHz of the HDPE and (a) black and (b) red PLA designs.	331
F.9	Diagram of the measurement set-up for the textile absorption analysis with the 300 GHz radar, using the focusing scheme from Fig. 5.4 and placing a plastic frame between the lenses and target to hold the textiles.	333
F.10	Spectrum comparison at 300 GHz of a reference metallic sheet and all clothing pieces inspected.	334
F.11	Diagram comparison of (a) purely specular reflection from a slightly-tilted metallic sheet with no clothing in front, and (b) the diffuse scattering created when placing the clothing.	334
F.12	Spectrum comparison at 300 GHz of a reference metallic sheet and all clothing pieces inspected, without using the focusing lenses.	335
F.13	Spectrum comparison of a reference metallic sheet and all fabrics considered, for) Configuration 3 of the 122 GHz radar.	335
F.14	Examples of non-resonant-based set-ups for permittivity and permeability extraction from a material's samples, where a VNA is used as the radiating source. (a) is the open-ended coaxial probe technique, (b) is a 2-port transmission and reflection method, and (c) is a 1-port reflection-only approach.	337
G.1	Block diagram of the 300 GHz imaging radar with numbered bandwidth steps (in red), and the bandwidths and frequencies currently used in the system (in light blue).	340
G.2	Proposed new microwave subsystem for the 300 GHz imaging radar using its DDS's image bands.	342
G.3	Proposed new RF-to-IF conversion for the 300 GHz imaging radar using two mixers.	344

List of Tables

1.1	Recent state-of-the-art active and passive imagers for security screening . . .	21
2.1	Relevant system parameters for the 300 GHz imaging radar	28
2.2	Comparison of real-time active imaging radars for security applications	34
2.3	GPU Specifications	36
2.4	CUDA signal processing configurations	43
2.5	Distribution of image time over kernels	45
2.6	Time per image in 100-image configurations.	52
2.7	CUDA signal processing configurations	53
2.8	Distribution of image time over kernels	54
2.9	Time per image in 100-image configurations.	57
2.10	Kernel timing comparison between the final GPU algorithm and a MATLAB-based implementation	58
2.11	Mean spot size for the framed spheres in Fig. 2.41	75
2.12	Effect of multistatic combination on image quality	78
3.1	Comparison among commercial MMIC solutions integrating full TRXs	93
3.2	Comparison among the two Silicon Radar GmbH MMICs used in this thesis .	95
3.3	Materials commonly present in security imaging scenarios	99
3.4	Mean error in mm for both MMICs with all measurement methods in Fig 3.15	104
3.5	Summary of current state-of-the-art high-frequency hybrid imaging systems .	115
3.6	System parameters for the developed hybrid imager	116
3.7	Tested imaging scenarios	148
4.1	Summary of the network architecture used, based on a modified AlexNet CNN	166
4.2	IR imaging datasets	182
5.1	Summary of recent imaging-based micro- and mm-wave NDT applications . .	200
5.2	Comparison of some of the current main labeling techniques	202
5.3	Summary of recent RF-based remote vital sign monitoring applications	218
F.1	Clothing pieces inspected with the 300 GHz and 122 GHz radars	333
F.2	Transmissivity losses for all fabrics considered at 300 GHz and 122 GHz . . .	336
G.1	Numbered steps in Fig. G.1 with currently used and maximum bandwidths (in GHz)	341
G.2	Summary of the proposed bandwidth improvements for the 300 GHz radar. .	344

Chapter 1

Introduction

1.1 Overview of the Millimeter-wave Spectrum

Electromagnetic radiation in the millimeter-wave (mm-wave) band extends from 30 GHz to 300 GHz [1], and it is called so because its corresponding wavelengths range between 10 mm and 1 mm. In the electromagnetic spectrum, this band is confined by the microwave (3 GHz to 30 GHz) and submm-wave (300 GHz to 3000 GHz) bands, as shown in Fig. 1.1. Multiple authors also include the THz band (usually bound by 100 GHz to 10 THz [2]) when discussing the mm-wave region, as intersection of the mm-wave, submm-wave, and low infrared (IR) bands, due to common radiation properties.

Figure 1.1: Region of the electromagnetic spectrum which includes the mm-wave, submm-wave, and THz bands.

1.1.1 History of the Millimeter-wave Spectrum

The use of mm-wave technology can be traced back to the 1890s, when wavelengths of 5 mm were first synthesized with spark-gap transmitters (TXs) [3]. In that same decade, submm-wave research began with the use of an IR heat source as a generator [4]. The nature of these two different approaches would be continued throughout the development of both frequency ranges; mm-wave solutions have been typically built around extensions of microwave technology, while submm-wave equipment has been mostly created through derivations of optical and IR sources.

In the decades up until World War II efforts were made to bridge the gap between optical and radiofrequency (RF) sources for mm-wave and submm-wave generation, but these were only partly successful. During the war, radar technology took on an important focus, but was

mostly developed around the microwave bands up to 24 GHz, slowing down research at higher frequencies [5]. After the war, mm-wave and submm-wave research resumed, driven mainly by molecular spectroscopy applications. In the mm-wave range, results were published with synthesized wavelengths as low as 1 mm by 1953 [6]. On the submm-wave side, spectrometers at 430 GHz were being developed from far-IR sources in 1950 [1]. Up to the early 1960s experiments in both bands were widespread, with developments in atmospheric propagation [7], radar systems [8], transmission lines [9], and high-power sources [10], but the development of the laser in the middle of the decade, once again, significantly halted research. Efforts resumed by the 1970s, when oscillators at 35 GHz started to become commercially available, thanks to the development of Gunn and IMPATT diodes [1].

With the development of the first affordable components in the 1970s, mm-wave technology started to be applied in diverse commercial applications, such as radiocommunication links at 60 GHz or automotive radars at up to 94 GHz [11]. Aside from commercial ventures, the space race between the United States of America and the Soviet Union motivated further progress during the Cold War. Submm-wave and mm-wave technology proved to be specially useful for spectroscopy in space applications, and numerous projects were launched by both countries [12]. This trend continued into the 1990s, when NASA and European organizations developed various ground- and space-based instruments [13, 14].

All throughout its development, technological progress in the mm-wave band has been notoriously delayed due to the so-called THz gap¹, unable to provide solutions between well-developed microwave electronics and mature optical technology for IR. Until the start of the new century, mm-wave applications were thus mostly limited to experimental and expensive projects, often in the radioastronomy and spectroscopy fields. Within the last two decades, continuous research and technological advancements have driven the development of affordable RF components for a wider variety of applications, ultimately reaching the introduction of commercial-off-the-shelf (COTS) products. Recently, the development of monolithic microwave integrated circuits (MMICs) at the mm-wave band has created availability of chip-sized RF systems at very low costs, forecasting even more commercial success for these frequencies [15].

1.1.2 Phenomenology of the Millimeter-wave Spectrum

Scientific interest in the mm-wave, submm-wave, and THz bands derives from the many properties that electromagnetic radiation possesses in these bands. One of the main properties is the high absorption values that these frequencies experience when encountering water molecules during propagation [16]. This can be verified with the reproduced measurements in Fig. 1.2. Although normally undesirable, these large attenuations are useful in various non-destructive testing (NDT), spectroscopy, security inspection, and healthcare applications, by allowing estimations of the water content in a radiated target. When dealing with live targets (such as humans or animals), radiation in these bands poses little health threat, as their photonic energy levels imply non-ionizing radiation (in contrast, for instance, to X-ray technology) [16].

While useful in some near-range applications, the absorption by water produces large attenuations during atmospheric propagation in long-range scenarios. This effect is specially harmful in these bands due to the relatively low-power radiation sources available [17], which do not allow compensating attenuation with an increase in transmitted power. In addition, absorption by water does not occur at a single frequency or fixed bandwidth, but rather at multiple resonant frequencies known as absorption lines [18], which create attenuation peaks with valleys between them (as seen in Fig. 1.3). This behavior creates the existence

¹A lack of efficient and practical components for signal generation and detection in the upper mm-wave and submm-wave regions [2].

Figure 1.2: Absorption coefficient of deionized water according to the signal's frequency, reproduced from [16].

of attenuation windows (i.e. bandwidths within the spectrum where attenuation is at a minimum) where mm-wave and submm-wave propagations are best designed. These windows allow relaxed operation at relatively-near distances (up to a few kilometers), but still become problematic for long-range propagation. The data in Fig. 1.3 also show increased attenuation with temperature, which is due to its effect on water vapor in the atmosphere. Other authors have also detected increase in attenuation depending on variables such as humidity levels, precipitation, or latitude [19,20]. On the other hand, when compared to transmissions in other regions of the spectrum (such as the visible or IR ranges), mm-wave signals perform better under low-visibility conditions due to nighttime, fog, dust, smoke, and other obscurants [21].

Figure 1.3: Atmospheric attenuation throughout the spectrum, with marked absorption peaks from oxygen and water molecules, reproduced from [19]. Above ~ 300 GHz windows become less broad as closer absorption lines appear and baseline attenuation increases.

The small wavelengths associated to these regions of the spectrum also allow for high-level integration in circuits thanks to component scaling, permitting the reduction in size of RF systems (ultimately trending to full system-on-chip designs in MMICs). Additionally, the diffraction limit for small wavelengths implies great angular resolution achievable with

small antenna apertures, allowing further reduction in a system's size [16]. In imaging applications, the spatial (or cross-range) resolution is what determines the sensor's capability for differentiating targets in close proximity, making it an extremely important benchmark. On the other hand, small wavelengths also produce very specific scattering patterns from reflecting targets. In rough surfaces where the micro-curvatures have comparable size to the transmitted wavelength, scattering is primarily diffuse (i.e. reflected back in a wide range of angles), while in smooth facets the reflections will be mainly specular (i.e. reflected in a single direction). In imaging applications, specular scattering can produce two detrimental effects, in the form of bright speckles and glints (i.e. localized powerful reflections when the reflected angle matches the antenna's aim), and vanishing detections (i.e. areas of the target where almost no power is returned back to the imager when the reflected angle does not match the antenna's aim) [22].

The other main property creating interest in these bands is their ability to penetrate a wide variety of non-metallic materials. This is specially useful in security screening applications where threats are concealed behind various fabrics, or in emergency response scenarios where a subject might be trapped under wood, plaster and other debris. As illustration, Fig. 1.4 shows absorption values through the spectrum for different textiles. It can be seen in the figure that in the mm-wave region, transmission will be only slightly affected by absorption, while as the frequency increases the attenuation will be more significant. In other measurements, the authors found that the 3-dB attenuation point for a variety of materials ranging from silk to wood was greater than 300 GHz [19]. In contrast, metallic and liquid materials are not likely to be penetrated. The conductive properties of metallic objects will cause most of the incident radiation to be reflected back, while the water content in liquid materials will produce very high absorptions. This contrast between penetrable and reflecting materials is useful in threat detection. In addition, blackbody radiation from a target's own body temperature presents emission peaks at frequencies such as 94 GHz where clothing transmissivity is also possible, making passive imaging a feasible application.

Figure 1.4: Absorption of various fabrics throughout the THz spectrum, modified from [23].

A beneficial side-effect of working at the mm-wave and submm-wave bands is the ability to operate with very large bandwidths. For communication applications this leads to higher data throughputs, while in frequency-modulated radars it allows for increased resolution in the range dimension. Additionally, occupying large bandwidths during transmission provides an additional layer of security against potential jamming techniques [24].

1.1.3 Applications in the Millimeter-wave Spectrum

Some of the current applications for the mm-wave band are summarized within the following fields:

- **Radioastronomy.** Astronomy and planetary science have been some of the main forces driving the development of mm-wave and THz technology [12]. Although atmospheric opacity and absorption lines limit large-distance propagation in this band, once the corresponding instruments are placed at adequate locations² large quantities of information can be extracted from cosmic emissions. Within this field, molecular line spectroscopy³ of gas and dust accumulations in star formations is the main focus. The data retrieved from this technique provide information regarding the chemical composition of matter in neighboring stars, planets, and other celestial bodies, allowing better understanding of the origin and growth of the universe. In addition, thermal emissions from blackbody radiation can also be found within the mm-wave and THz bands, enabling another source of cosmic information. Lastly, according to Hubble's Law, emissions in the visible spectrum from far-off formations are redshifted into lower frequency bands (including the THz range) allowing measurements of their movement. Even though the use of mm-wave and THz technology for radioastronomy has been present since the 1960s [25], many recent instruments and large-scale projects in this band have been successfully developed. These include the Atacama Large Millimeter/submillimeter Array⁴ in the Chilean desert [26], the space-based Herschel Space Observatory⁵ [27], the Rosetta space probe⁶ [28], and the Planck spacecraft⁷ [29], among other notable projects⁸.
- **Spectroscopy.** For Earth-based spectroscopy, plasma and gas studies have been the two principal areas of research [18]. In plasma studies, mm-wave and THz radiations can be used to infer electron density [35], temperature distribution [36], and even 2-D or 3-D images with varying penetration [37]. In this field, mm-wave instruments have proven to be apt for operation in harsh environments, which make them specially suitable for the heat conditions around tokamak-confined plasma. As for the molecular analysis of gases, the advantage of working in the mm-wave and THz band comes from the exponential increase of some absorption lines with frequency. The identification of gas signatures in these bands can be used to monitor polluting emissions [38], remotely detect specific chemicals and aerosols [39], or classify explosives [40], among other uses.
- **Weather and Climate Studies.** The spectroscopy and astronomy fields have also overlapped to develop uses for the mm-wave and submm-wave regions regarding weather forecasting and climate analysis. In this area of study, the atmospheric absorption at mm-wave frequencies can be used to extract vertical line measurements of the atmosphere's humidity and temperature from ground-based or satellite sensors [41], as well as for other applications such as cloud profiling [42], cloud ice measurement [43], or precipitation analysis [44]. On the other hand, if these absorption windows are

²Usually dry and high mountaintops, flying aircrafts, or orbiting stations, to avoid as much atmospheric propagation as possible.

³Detecting and classifying emission and absorption signatures from gases.

⁴An array of 66 antennas operating from 31 GHz to 1 THz to study star births, fully operational since 2013.

⁵Using equipment from 480 GHz to the IR region to analyze galaxy and star formation from 2009 to 2013.

⁶A spacecraft for comet observation launched in 2004 and operational until 2016, with equipment at 190 GHz and 562 GHz, among other frequencies.

⁷A space observatory operated from 2009 to 2013 to study cosmic microwave background radiation between 30 GHz and 857 GHz

⁸Such as the Combined Array for Research in Millimeter-wave Astronomy [30], the MetOp satellites [31], the Stratospheric Observatory for Infrared Astronomy [32], the Superconducting Submillimeter-wave Limb-Emission Sounder [33], and the Submillimeter Array [34].

avoided, remote sensing of the Earth’s surface can be performed from orbiting platforms. Gas spectroscopy analysis can also be extended to the Earth’s atmosphere to measure ozone levels [45], greenhouse effects [46], and other climate change processes. Additional geoscience applications include land-based snow height measurements [47], oil spill monitoring [48], and oceanic wind estimations [49], among others.

- **Security Imaging.** One relatively-recent research area for mm-wave and low-THz radiation has been the use of RF imaging for various security applications [20]. Thanks to some of the properties of radiation in these bands, electromagnetic penetration through a large variety of materials is possible, allowing non-invasive inspection through containers, luggage, and even person-borne clothing. The transmission through these obstacles enables the detection of concealed items, such as weapons, drugs, and contraband. For this task, mm-wave radiation poses several advantages versus traditional scanning technologies (such as X-rays) due to their non-ionizing levels and farther detection ranges. This research has of course driven the development of whole-body imagers for passenger control in locations such as airports or train stations [50], but also for other settings such as post offices (where imaging can help scan incoming mail [51]). The current worldwide security threats from terrorism have driven mm-wave security imaging into one of the few commercially-feasible applications of the frequency band, with several companies producing COTS imaging products.
- **Communication Systems.** Another area that has recently been motivating the development of mm-wave and submm-wave technology deals with communications systems. In this regard, most designs for future 5G mobile links propose extension into the mm-wave spectrum, which provides the large bandwidths needed to accommodate the desired data throughputs [52]. In addition to cellular networks, the IEEE 802.11ad⁹ standard for WiFi communication has been built around a physical layer at 60 GHz [53]. Aside from civilian applications, mm-wave communication systems have previously been proposed in many military settings, such as soldier-to-soldier communications [54], armored vehicle systems [55], or satellite links [56].
- **Material Characterization.** The imaging capabilities of mm-wave and submm-wave technology have also been exploited for applications apart from security screening. In this regard material evaluation and NDT are some of the most prominent uses [57]. For these tasks, the ability to penetrate diverse dielectrics in these frequency bands is used to remotely inspect materials such as cork [58], plastics [59], coal [60], and concrete [61], among others. The resulting images can be used to detect superficial and hidden defects, making them extremely useful for quality control processes in the pharmaceutical [62], construction [63], and food industries [64]. When combined with spectroscopy techniques, NDT results can also help extract specific material properties like density, surface roughness, and dielectric constants, ultimately allowing classifying polymers, alloys, and different materials.
- **Healthcare.** The very small wavelengths encountered in the mm-wave and submm-wave bands allow long-distance detection of μm -level displacements in static objects. This property has been used since the 1970s to remotely monitor a subject’s chest movements due to respiration and heartbeat [65]. Since then, vital sign monitoring has evolved to overcome challenges such as multi-subject monitoring, body movement cancellation, and fall detection, among others [66]. Additionally, the water-based absorption in these bands and their imaging capabilities have been used to scan breast cancers [67], burn injuries [68], human corneas [69], and other tissues [16].

⁹The newest standard at the time of writing, without including the IEEE 802.11aj standard (which also uses a mm-wave center frequency, at 45 GHz, but is mostly used in China).

- **Autonomous Driving.** The most recent incentives for mm-wave research have come from the automobile industry. Vehicle manufacturers are currently trending towards autonomous driving solutions, which require sensors for precise localization and navigation [70]. The ranging capabilities of radars in the mm-wave and THz bands allow for detection of targets up to hundreds of meters with great resolution, making them a prime candidate for vehicle-, obstacle- or pedestrian-detection. In this field, mm-wave instruments can be used for many applications, like radars for collision avoidance [71], road curvature estimation [72], and terrain sensing [73], as well as many others.
- **Other applications.** Aside from the various research fields discussed, mm-wave and submm-wave technology have been successfully used in many other areas. Recent notable uses of these bands have been, for example, space debris detection [74], gesture sensing [75], aircraft landing assistance in low visibility [76], crop and vegetation monitoring [77], foreign object detection on runways [78], cavity profiling in underground mines [79], and speech detection [80].

Within these applications, mm-wave security imaging has been the main focus during the development of the radar systems discussed in this thesis in Chapters 2 and 3. Additional experiments related to the healthcare, material characterization, and data acquisition fields are presented in Chapter 5.

1.1.4 Outlook and Trends

For the mm-wave band, in a 1995 review [11], automotive radars were already expected to be one of the main driving forces behind future commercial success. This has partially proven to be true with the rise of autonomous driving research in the last few years. On the other hand, the expectations for the telecommunications industry to be the other main actor in millimeter-wave expansion have not fully materialized. Nonetheless, it is true that recent developments¹⁰ indicate that widespread mm-wave devices will soon become a reality.

A somewhat sudden push in mm-wave development has come from security screening applications. Due to global terrorist threats in the past two decades, this field has gathered great interest, and has in turn focused on the mm-wave band for new technological solutions. In a 2017 review, an ongoing seminal conference dedicated exclusively to mm-wave imaging reported more than 466 publications over 20 years, with the main frequencies discussed being 35 GHz and 94 GHz [82]. A 2014 report by the United States Department of Homeland Security (DHS) already encountered up to 11 different manufacturers of passive mm-wave scanners, with prices ranging from \$19900 to \$200000, indicating a fairly wide market for such a novel application [83]. Not surprisingly, a 2012 report for the United States Congress found that the percentage of full-body scanners in their airports which used X-ray technology had decreased from 44% to 35% in favor of mm-wave solutions, in just 5 months [84]. On the other hand, a 2010 survey cited in the congressional report showed that 44% of scanned passengers were worried about privacy issues, 27% were concerned with potential health risks, and 20% questioned potential added delays. Additionally, operators showed concern regarding the effectiveness of scanners against children or passengers with disabilities. At a higher level, government agencies in charge of airport security introduced automatic threat recognition software in their scanners in 2011, with the intention of completely eliminating human operators in the near future. At a technological level, these demands require several improvements in mm-wave imaging, specially in acquisition times, reliability, and cost. Expected future trends to meet these requirements are the introduction of multimodal systems, innovations in RF integration, and increased research in non-mechanical beam steering [82].

¹⁰Such as official allocation of several GHz of mm-wave bandwidth for 5G communications in the United States in 2016 [81].

As for the submm-wave and THz bands, a 2014 study summarizes some of the statistics of the growth in this region of the spectrum [85]. The Internet query rate for terms related to “THz” have been in decrease since 2004. Within these searches, “THz imaging” and “THz spectroscopy” are still the most searched terms. Although queries are decreasing, scientific publications and patents with these terms present exponential growth since 1990. Altogether, these findings seem to suggest that these frequency bands have not yet gathered significant public impact, but are still scientifically interesting. This is further proven by the linear growth in THz system manufacturers since 2004, with most of them grouped in North America, Germany, and Japan. Within global manufacturers, 61% view imaging as one of their main applications, while 52% and 39% do the same for security inspection and NDT, respectively. The commercial growth in these bands can be demonstrated by the 60% reduction in cost for full THz systems between 2002 and 2013, but the authors determine that demand for these instruments is still lower than the worldwide offer (which in part creates the still-high costs common in these bands). The authors finish by concluding that scientific activity in the THz field has been continuously increasing, and this trend is expected to continue. A 2017 review of scientific activity in the submm-wave and THz bands states the current challenges of this region of the spectrum when applied to various tasks [86]. The common obstacle for imaging, medical applications, security inspection, and space technology is the high cost of components in these bands. Another shared challenge is the usually-large size of these components and systems, specially detrimental for the imaging and space fields. One last common obstacle is the acquisition speed at which these systems are capable of refreshing their results, most notably in NDT and security inspection. The authors conclude that overcoming these obstacles is necessary in order to reach this spectral region’s full commercial and operational potential.

1.2 Imaging in the Millimeter-wave Spectrum

As already mentioned, imaging is one of the main applications for technology in the mm-wave, submm-wave, and THz bands. Compared against optical and IR imaging solutions, the ability to penetrate various obstacles such as dust, smoke, or clothing is the main attraction of working in lower frequencies. When compared against microwave technology, mm-wave and submm-wave systems offer increased resolution with the same antenna sizes due to the decrease in wavelength. Within all possible sub-applications of mm-wave and submm-wave imaging, this section will be focused on security screening, as it has been the main focus of the radar systems developed during this thesis.

The first subsection deals with the two main types of imaging encountered in these bands. The following subsection details the different hardware and signal processing architectures commonly employed in these systems. Finally, the last subsection regarding imaging is a summary of some of the state-of-the-art systems in commercial and academic references.

1.2.1 Comparison between Active and Passive Imaging

The imaging sensors used can typically be classified into two categories: active and passive. In active imaging, the processed signals are dominated by the scene’s scattered reflections of radiation transmitted from an RF device, thus providing information about a target’s reflectivity. Passive imaging, on the other hand, makes use of the scene’s self-emitted thermal radiation in specific frequency bands [87] (providing information about the target’s emissivity), as well as the scene’s reflections from surrounding uncontrolled sources (in relation to the target’s reflectivity) which are usually of atmospheric origin.

For both types of imaging received power ($P_{received}$) can be modeled [88], in a simplified manner, by the following equation

$$P_{received} = P_{thermal} \cdot \varepsilon + P_{surroundings} \cdot r + P_{background} \cdot t \quad (1.1)$$

where $P_{thermal}$ is the target's own radiated power due to its temperature (which is emitted according to its emissivity, ε), $P_{surroundings}$ is the radiated power from the target's surroundings (which reach the imager after reflecting off the target by a factor known as reflectivity, r), and $P_{background}$ is the radiated power from the area directly behind the target (which passes through the target depending on its transmissivity, t). The three ratios (ε , r , and t) are all magnitudes between 0 and 1, where $\varepsilon = 1$ defines a perfect emitter (also known as a blackbody), $r = 1$ denotes an ideal reflector, and $t = 1$ indicates a completely transparent material.

In active imaging, $P_{surroundings}$ can be decomposed into various contributions

$$P_{surroundings} = P_{transmitter} + P_{multipath} + P_{clutter} \quad (1.2)$$

where $P_{transmitter}$ is a fraction of intentionally-transmitted power that reaches the target in a direct manner, $P_{multipath}$ is the fraction that reaches the target after multiple reflections in the scene, and $P_{clutter}$ comes from the naturally-occurring radiation in the scene (such as atmospheric emissions). Although, $P_{transmitter}$ is attenuated (by a factor α mainly dependent on atmospheric conditions) with respect to the actual transmitted power, its value is still much larger than thermal emissions from the scene or the target, and the sensor usually lacks enough dynamic range to detect both transmissions. Thus, in active imaging

$$P_{received} \approx P_{transmitted} \cdot \alpha \cdot r \quad (1.3)$$

On the other hand, in passive imaging $P_{transmitted} = 0$. In this case, radiation from thermal origin is the dominating contributor to received power. Power received from an ideal thermal emitter can be related to its temperature through a magnitude called spectral brightness (B_f) [19], for which Planck's Law states

$$B_f = \frac{2hf^3}{c^2 e^{\frac{hf}{kT}} - 1} \quad (1.4)$$

where h and k are Planck's and Boltzmann's constants (respectively), c is the speed of light in vacuum, f is the operational frequency, and T the object's temperature. In mm-wave frequencies, Planck's Law can be approximated by the Rayleigh-Jeans Law which states

$$B_f \approx \frac{2kTf^2}{c^2} \quad (1.5)$$

indicating a direct relationship between object temperature and spectral brightness. Factoring in a non-ideal emitter (modelled by its emissivity), the brightness becomes

$$B_f \approx \frac{\varepsilon 2kTf^2}{c^2} \quad (1.6)$$

and when the detector uses an antenna aperture A_r , the received power can be computed by integrating B_f over A_r and the solid angle (Ω) associated to the object at the receiver's (RX's) location. This is summarized by

$$P_{thermal} = \int_{A_r} \int_{\Omega} B_f \cos(\theta) d\Omega dA_r \quad (1.7)$$

In the case the RX operates with a bandwidth (B) instead of at a single frequency, Bf would also need to be integrated over B . Through this formula, potentially, even objects

at the same temperature could be distinguished due to their different emissivities. Such behavior is useful, for instance, when trying to detect a concealed metallic threat which has reached temperature equilibrium with the skin it presses against. The very different emissivities for both materials ($\varepsilon \approx 0$ for metal and $\varepsilon \approx 0.9$ for skin, at 500 GHz [20]) would create large power contrast. Additionally, the metallic object's low transmissivity ($t \approx 0$, at 500 GHz [20]) would block emitted radiation from the skin, creating a noticeable “hole” in the resulting image.

The different sources of radiated power in both types of imaging are illustrated in Fig. 1.5.

Figure 1.5: (a) Active imaging scenario, where received power is dominated by reflections of the transmitted signal. (b) Passive imaging scenario, where received power presents approximately equally-powerful contributions from background, clutter, and the subject's own emissions.

Due to their unsimilar principles of operation, each class of sensor has different advantages. Active sensors usually present a larger dynamic range and stronger reflections from metallic objects, while passive devices¹¹ eliminate the need for powerful signal sources (reducing cost, architectural complexity, and probability of interception) and can present lower scattering attenuation due to one-way propagation. In consequence, the disadvantages of these sensors also differ: in active imaging, specular reflection is usually the primary problem [22], while detector cooling and background separation are the typical obstacles of passive imaging.

Because they rely partly on large temperature contrasts between surroundings and targets to establish their dynamic range, passive imagers commonly present better performance in outdoor environments, where the sky's low mm-wave temperature (~ 100 K in clear conditions [90]) provides a great background. The sky also serves as a great illuminator for highly-reflective threats (such as metallic weapons), which will then appear as cold objects in the resulting image, contrasting highly with the subject's warm body (which has driven researches to use incoherent global illuminators to mimic sky radiation indoors [89]). Additionally, outdoor scenarios usually include less clutter (in the form of other reflective objects or powerful thermal emitters), which increases an image's quality.

These considerations are illustrated in Fig. 1.6, where outdoors and indoors passive imaging is compared. In the outdoors scenario, the gun's high reflectivity and low emissivity produce a received temperature equal to the cold sky's, while the human's high emissivity and low reflectivity create the inverse effect (where the human's received temperature is almost the same as his actual body temperature). The large difference between received gun and human temperatures creates good contrast for the sensor to generate a high-quality image. On the other hand, in the indoors scenario the gun reflects clutter temperature (which includes the warm room's atmosphere) which reduces contrast.

¹¹Sometimes also called radiometers.

Figure 1.6: (a) Passive outdoors imaging scenario, where received temperature from threat and subject presents a good contrast thanks to the sky's cold radiation. (b) Passive indoors imaging scenario, where received temperature from threat and subject presents a low contrast due to the room's warm surroundings.

Furthermore, passive detectors require very good sensitivities, because blackbody radiation in the operational bands is noise-like and limited to picowatt levels [19]. In this regard, both types of imaging are limited by atmospheric attenuation at the operational frequency bands to detection ranges of up to a few meters, but the limit is more strict in passive scenarios (where received radiation is weaker). This harder limit for passive systems also reduces their operability, when compared against active imagers, under atmospheric obscurants and person-worn clothing.

Although originally lagging behind active technology development (due to an initial lack of high-sensitivity detectors [90]), a large number of commercial security imagers are passive systems, partly due to their simpler architecture. The additional complexity in active architectures comes from the need to include power sources and TXs, which at the considered frequency bands usually require multiples stages of signal conditioning. This is due to the lack of signal sources directly in the mm-wave and submm-wave spectrums, which forces signal generation to come from the microwave or IR realms. In the case of microwave sources, mm-wave and submm-wave signals are generated through upconversion with frequency multipliers, while for IR sources photomixing for downconversion from lasers is a common design.

Current imaging technology for concealed weapon detection in the mm-wave and THz ranges is traditionally limited to a single type of sensor. Nevertheless, no individual threat detection method is perfect under all imaging conditions, and undetected cases in security applications can imply great risks [19]. Due to these circumstances, multimodal imaging has recently began to gather research attention. The simultaneous use of multiple frequency bands or sensor types provides complementary information about the scene, which can later be merged efficiently with sensor fusion techniques. Multifrequency examples of single-type sensors are common [91], and recent research around Europe has been focused on developing combined active and passive imagers (in some cases also including a support sensor in the visible range) [92, 93].

Regarding the near future, in active imaging, transistor technology is currently starting to overtake Schottky diode solutions, while large bolometer arrays are emerging as candidates for standard passive configurations. Moving further into the future, both types of imaging

could benefit from further development in MMIC technology [86]. Full RXs in chips, which imply lower costs and size, would allow creation of large imaging arrays to increase acquisition speeds. For efficient MMIC solutions, packaging technology still needs to progress [94].

1.2.2 Imager Architecture

Within the two types of imaging, the sensors used can be very varied in their complexity. Passive imaging can range from a simple detector at a specific frequency to an array of coherent sensors each covering a channel of a wide bandwidth. Similarly, active imaging can be done simply with an incoherent illumination source and a broadband detector, or with more elaborate architectures using distributed coherent transceivers (TRXs) emitting modulated waveforms. Because the two main systems described in this thesis are active imagers, this section will be focused on active sensor architecture only.

Active imagers in the mm-wave and submm-wave bands are radar systems. A radar by definition is an instrument which allows measuring one or several magnitudes associated to a target (such as range, speed, or reflectivity) by correlating a transmitted electromagnetic signal with a delayed reflection. Although radar development has been historically associated to pulsed systems (where target range can be simply measured as the time-of-flight of the returned pulse), most active imagers transmit wide bandwidths in frequency-modulated continuous waveforms to exploit the ranging capabilities and resolution that these signals offer. Within these systems, frequency-modulated continuous wave (FMCW) radars are the common choice for waveform selection.

Instead of pulsing their transmission, FMCW radars transmit and receive simultaneously. The transmitted waveform is a frequency modulation of a continuous sinusoidal signal, of which the most simple and common is a linear frequency sweep (also called chirp). An ideal target would then simply return the same modulated signal, delayed and attenuated. In the RX, the reflected signal is continuously mixed with the transmission producing an intermediate frequency (IF) signal whose frequency is proportional to the delay (and thus range). With simply a Fast Fourier Transform (FFT) the time-domain IF signal can be converted into a spectrum which directly illustrates a range profile.

In general, FMCW radars present several advantages when compared to pulsed systems, such as:

- **Low levels of peak transmit power.** Because pulsed systems transmit for only small fractions of time, their peak transmitted power is usually large in order to effectively reach their target. On the other hand, since transmitted energy is the product of peak transmitted power times signal duration, FMCW systems can achieve the same transmitted energy as pulsed systems but with lower peaks, because they transmit in a continuous manner. Lower power peaks allow FMCW systems to remain undetected with more ease, interfere less with other TXs, and use less powerful energy sources.
- **Measurements in close range.** Since pulsed systems do not allow reception during transmission, any returned signals during the transmit pulse will not be processed. When converting the pulse's duration into range, this effect creates a minimum detectable distance for targets. FMCW radars on the other hand are continuously receiving reflections, and are thus perfect for close-range applications such as security imaging.
- **Range resolution.** The frequency modulation of radars is dependent on transmitted bandwidth, which is readily available in the mm-wave and submm-wave bands. To increase their bandwidth, and thus resolution, simple pulsed systems reduce their pulse duration, which demands faster switching and higher peak power (which are not so

easily attainable at high frequencies). This relationship between pulse duration and bandwidth in simple pulsed radars is also a disadvantage, because in FMCW radars chirp duration and bandwidth are independent variables¹².

However, FMCW radars also come with their own disadvantages. For example, because these systems use their signal's frequency for range information, a Doppler shift in a target will erroneously spread that target's range profile¹³. Because FMCW radars transmit and receive simultaneously and their power levels are low, any magnitude of TX leakage can severely mask close targets, which usually forces the use of more expensive and complex two-antenna configurations. Lastly, the linearity of the frequency sweep is of utmost importance, and any deviation from a perfectly linear sweep will reduce the effective range resolution and introduce parasitic modulations in IF.

Fig. 1.7 presents a block diagram of a generic FMCW architecture. A chirp generator produces a frequency sweep at the microwave band, which is later upconverted by passing through a multiplier chain. Because mm-wave multipliers usually present low efficiencies, it is necessary to amplify the upconverted chirp before transmitting. A two-antenna configuration provides simultaneous transmission and reception with good isolation. In the RX channel, an amplifier is the first stage, as is common in most RF systems. Afterwards, a mixer generates the IF signal by multiplying a sample of the transmitter chirp with the received reflections. The IF signal is then low-pass filtered to eliminate unwanted harmonics and spurs, while maintaining the desired beat frequencies. Finally, the IF signal is sampled with an analog-to-digital converter (ADC), which can store or transfer the signal samples for processing.

Figure 1.7: Block diagram of a generic single-antenna FMCW architecture.

Assuming a linear frequency sweep described by

$$f(t) = f_0 + \frac{BW \cdot t}{T_m} \quad (1.8)$$

where f_0 is the initial frequency, BW is the transmitted bandwidth, and T_m the duration of each chirp, the ideal transmitted signal $[x_{TX}(t)]$ in FMCW operation is a sinusoid with constant amplitude (A) and modulated frequency defined by

$$x_{TX}(t) = A \cos(2\pi(f_0 t + \frac{\gamma t^2}{2}) + \varphi_0) \quad (1.9)$$

¹²Chirp duration must still be long enough for the reflected signal to significantly overlap in time with the transmission, or else the IF signal will be useless.

¹³This can be nonetheless corrected simply with comparisons over repeated and subsequent chirp signals from the same target, or by transmitting specific waveforms.

where φ_0 is the initial phase, and γ is known as the chirp rate and is a substitution for

$$\gamma = \frac{BW}{T_m} \quad (1.10)$$

The formula in (1.9) only describes the first chirp in the sequence. For the n^{th} chirp, the time variable can be replaced by

$$t = nT_m + t_s \quad (1.11)$$

where t_s is the time since the start of the n^{th} chirp. With this in mind, the n^{th} chirp is defined by

$$x_{TX}(t) = A \cos(2\pi(f_0(nT_m + t_s) + \frac{\gamma t_s^2}{2}) + \varphi_0) \quad (1.12)$$

An ideal target will simply reflect a time-delayed and attenuated copy of the transmitted wave. The delay from two-way propagation (Δt) will be dependent on the range to the target (R) by

$$\Delta t = \frac{2R}{c} \quad (1.13)$$

where c is the speed of light in vacuum. Applying this delay and the attenuation to (1.14), the received signal $[x_{RX}(t)]$ can be defined by

$$x_{RX}(t) = B \cos(2\pi(f_0(nT_m + t_s - \Delta t) + \frac{\gamma(t_s - \Delta t)^2}{2}) + \varphi_0) \quad (1.14)$$

where $B < A$. In the RX's mixer, the transmitted and received sinusoids will be multiplied. Knowing that

$$\cos(\alpha_1)\cos(\alpha_2) = \frac{\cos(\alpha_1 + \alpha_2) + \cos(\alpha_1 - \alpha_2)}{2} \quad (1.15)$$

and that only the difference term will remain unfiltered, the mixed signal $[x_{MX}(t)]$ can be described as

$$x_{MX}(t) = \frac{AB}{2} \cdot \cos(2\pi(f_0\Delta t + \gamma t_s\Delta t - \frac{\gamma(\Delta t)^2}{2})) \quad (1.16)$$

which is the equation of a sinusoid signal $[C \cos(\varphi(t))]$ whose instantaneous frequency $[f_i(t)]$ is

$$f_i(t) = \frac{\partial \varphi(t)}{2\pi \partial t} = \gamma \Delta t = \frac{2R \cdot BW}{c \cdot T_m} \quad (1.17)$$

indicating that after mixing, the IF signal in an FMCW radar will have a frequency dependent on the range to the target. This frequency (f_b) is commonly called the beat frequency, and can also be illustrated with Fig. 1.8. In this figure, taking place in the RX's mixer, the blue sawtooth signal is a sample of the transmitted wave, while the red trace is the reflected signal.

Noting the similarity between the triangles defined by $(f_b, \Delta t)$ and (BW, T_m) , f_b can be calculated as

$$f_b = \frac{\Delta t \cdot BW}{T_m} = \frac{2R \cdot BW}{c \cdot T_m} \quad (1.18)$$

Figure 1.8: Time-frequency diagram of sawtooth FMCW operation.

The relationship between f_b , T_m , and BW is illustrated in Fig. 1.9. First, Fig. 1.9(a) shows the evolution of the beat frequency according to target's distance when the bandwidth is set to 10 GHz for three different chirp times. Conversely, Fig. 1.9(b) presents the beat frequency's progression for three different bandwidths when the chirp time is 1 ms. All orders of magnitude considered are common for state-of-the-art imagers¹⁴. As expected, wider bandwidths and faster chirp times imply higher beat frequencies. Because these beat frequencies must be sampled for processing, this relationship poses a tradeoff between, on one side, image resolution (from bandwidth) and acquisition time (from chirp time), and on the other, baseband limitations (from sampling speed, data throughput, and file sizes).

Figure 1.9: (a) Beat frequency as a function of the range to the target and chirp time, when the transmitted bandwidth is 10 GHz. (b) Beat frequency as a function of the range to the target and bandwidth, when the chirp time is 1 ms.

As previously mentioned, an FFT is typically used to extract f_b from the mixer's time-domain output. The bandwidth of each FFT bin can be then identified as the f_b resolution (Δf_b), and will be approximately equal to the inverse of T_m times a factor (α) corresponding to the type of signal windowing used [95], as represented by

$$\Delta f_b \approx \frac{\alpha}{T_m} \quad (1.19)$$

¹⁴For example, the two systems presented in this thesis transmit a 150 μ s chirp with 27 GHz and 15 GHz of bandwidth, respectively.

Substituting (1.19) into (1.18), the range resolution (ΔR) in an FMCW radar is determined by

$$\Delta R = \frac{c \cdot \alpha}{2 \cdot BW} \quad (1.20)$$

The formula in (1.20) is illustrated in Fig. 1.10(a), for a scenario where $\alpha = 1$. For a 1 cm range resolution, which is commonly considered the necessary minimum to accurately represent smaller threats and separate various clothing layers, a bandwidth of at least 15 GHz is needed. However, this corresponds to the scenario where a $\alpha = 1$ from using a rectangular window¹⁵, so an even wider bandwidth would be necessary in a practical case where $\alpha > 1$. The actual bandwidth necessary for a range resolution of 1 cm when $\alpha > 1$ can be seen in Fig. 1.10(b).

Figure 1.10: (a) Range resolution as a function of the transmitted bandwidth, when $\alpha = 1$. (b) Necessary bandwidth for a range resolution of 1 cm when $\alpha > 1$.

In a traditional radar system the signal-to-noise ratio (SNR) at the RX's output (SNR_o) can be computed as the ratio of signal and noise powers as

$$SNR_o = \frac{\alpha P_t G_t A_{e_r} \sigma}{k T_e B L_{sys}} \quad (1.21)$$

where α is the atmospheric attenuation, P_t is the transmitted power, G_t is the transmit antenna gain, A_{e_r} is the effective aperture of the receiving antenna, σ is the radar cross section (RCS), R is the range to the target, k is Boltzmann's constant, T_e is the RX's equivalent noise temperature, B is the RX's bandwidth (which with FFT processing will be equal to the FFT binwidth, B_{FFT}), and L_{sys} are system losses.

Knowing that the effective aperture can also be defined by

$$A_{e_r} = G_r \frac{\lambda^2}{4\pi} \quad (1.22)$$

where G_r is the receiving antenna's gain, the formula in (1.21) can also be written as

$$SNR_o = \frac{\alpha P_t G_t G_r \lambda^2 \sigma}{(4\pi)^3 R^4 k T_e B L_{sys}} \quad (1.23)$$

¹⁵Rectangular windows are not usually chosen in security screening because of their powerful side-lobes, which could mask the less powerful range layers nearby.

which assumes a transmitting mechanism which does not confine its radiation into a single spot, but rather distributes it in a sphere. For imaging radars with a focusing system, the illuminated spot in the target can be thought of as a point-like radiating source of αP_t power, which eliminates the $\frac{1}{4\pi R^2}$, σ , and G_t dependencies from (1.21)¹⁶. This substitution concludes with an SNR for a focused imager defined as

$$SNR_o = \frac{\alpha P_t G_r \lambda^2}{(4\pi R)^2 k T_e B L_{sys}} \quad (1.24)$$

Hardware and Scanning Configurations

A high-level classification of imagers can be made regarding the way they image the scene. In this classification, three main types of imagers can be defined: handheld, portals, and standoff. A handheld device will be a compact and battery-powered sensor that can be carried by an operator throughout the scene. The advantages to this class are its flexibility and relative low cost, while its drawbacks are small range to the target and sensitivity. Portals are the most commonly-seen class in commercial imagers deployed in airports and other locations. In this architecture, the subject stands inside an arc or circular enclosure and is imaged from multiple angles. The main advantage to this imager type is the ability to scan the subject from various viewpoints simultaneously, while its primary disadvantage would be the need to ensure that the subject remains static and nearby in a pre-determined pose. Lastly, standoff imagers scan the scene from a large range¹⁷, and can often do so while the target is moving and without consent. The main appeal of this class is the ability to perform imaging of noncompliant subjects at a safe distance, while its main disadvantages are usually cost and lack of multiple viewpoints. For future security checkpoints the trend is to develop walkthrough corridors (where one or more TRXs are placed) that would allow imaging of subjects without stopping their flow, increasing passenger throughput, as requested by major stakeholders like the DHS [98] or the Transport Security Administration (TSA) [99]. All these classes are illustrated in Fig. 1.11.

Figure 1.11: (a) Handheld W-band sensor, reproduced from [100]. (b) The commercial ProVision2 portal imager from L3 Security & Detection Systems, reproduced from [101]. (c) 300 GHz standoff imager, reproduced from [102].

For signal generation, it has already been mentioned that the mm-wave and submm-wave bands are severely limited by the THz gap's lack of sources. Although some imagers

¹⁶Valid only for a metallic and spherical target with an RCS larger than the spot size, in order to intercept and reflect all incident power [103].

¹⁷At least 5 m are considered necessary for survival when an explosive device is involved [20].

using photomixing to downconvert IR lasers and diodes exist, most current imagers generate their transmissions from upconversion of the microwave range. In these lower bands, several signal sources can be used, divided into two main categories: vacuum technology and solid-state devices. Vacuum tubes, such as the gyrotron, magnetron, or klystron, are high-power devices capable of producing megawatts of peak output. However, these instruments are usually expensive, bulky, and require restrictive power supplies and packaging conditions. On the other hand, solid-state technology, such as Gunn diodes, IMPATT diodes, or various transistors, are low-power solutions that allow low costs and small sizes. Due to the relaxed power requirements of FMCW radars, these systems commonly use solid-state sources in the microwave region, later upconverted with cascading Schottky diode multipliers (which are becoming COTS devices with great efficiencies [96]). Within solid-state sources, direct digital synthesizers (DDSs), voltage-controlled oscillators (VCOs), and dielectric resonant oscillators (DROs) are some of the commonly-used devices for chirp generation. Among these, DDSs have solidified themselves as the main source in recent imagers due to their many advantages, such as low cost, good linearity, fast switching, and flexible implementation [97].

For the RX side of the imager, several architectures are available. The simplest method would be direct detection, which at mm-wave frequencies is usually done with a simple GaAs Schottky diode, while at the submm-wave and THz bands bolometers are also used. A direct detection scheme trades subpar sensitivity for low cost and reduced complexity. A more selective approach would be a tuned RX using low-noise amplifiers (LNAs) before the detector, to filter unwanted frequencies and amplify the desired signal. This technique is sometimes not implemented due to difficulties of manufacturing LNAs at such high frequencies. The most used RX architecture is the heterodyne configuration, which uses a mixer to downconvert the reflected frequencies into IF, where filtering and amplification are more manageable. The main downside to this scheme is the existence of conversion losses in the mixer, which reduces its efficiency when compared to a simple detector.

Aside from the RF components necessary for signal generation and reception, an imager must also include a radiating element to transmit. The reduced dimensions of antennas at mm-wave frequencies do not provide the necessary directivity to concentrate the transmitted signal into small areas, which is desired in order to image the field of view (FoV) with enough resolution. This is specially detrimental in standoff systems (where range to the target is farther), because the area where the antenna is concentrating its radiation grows with distance. To solve this issue, imagers usually include additional components such as large reflectors or multi-TX arrays to provide enhanced focusing capabilities. Furthermore, thanks to the proximity to the IR and visible spectrums, quasioptical elements (such as beamsplitters, mirrors, or lenses) can be jointly used with traditional radiating components.

To generate an image, the focused radiation spot must be displaced for each pixel within the FoV, for which a scanning mechanism is needed. Two main scanning techniques are usually encountered in imagers: mechanical and electronic. Mechanical scanning includes motors in the reflector system (typically acting on a mirror or reflector) to shift the orientation of the focusing spot in a pre-determined pattern covering the entire FoV. On the other hand, electronic scanning is typically implemented with the use of a phased array, which allows beamforming techniques to steer the spot without physically moving any component¹⁸. Fully-electronic active arrays in mm-wave bands are still costly¹⁹, plus they require additional post-processing calibration to compensate the different gains and offsets of each element, so most imagers still use some type of mechanical scanning to cover the FoV. On the other hand, a drawback of mechanical scanning compared to fully-electronic scanning is the sequential nature of the positioning for each measurement, which reduces imaging speed. Combinations of these two methods have also been reported, such as the use of focal plane arrays (FPAs) of

¹⁸The individual phase and amplitude of the signal in each radiating element in the array can be modified to create constructive/destructive interference patterns in a selected direction.

¹⁹Somewhat in contrast to passive imagers, where arrays of thousands of elements are common.

TRXs where an array is placed in the focal plane of a reflector or lens and each element scans only part of the FoV in a mechanical approach. Fig. 1.12 illustrates the different scanning approaches mentioned.

Figure 1.12: (a) Single TRX scanning where the reflector must be mechanically moved to illuminate each spot. (b) Multi-TRX array scanning where beamforming techniques are used to illuminate each spot without need for movement. (c) Combination of mechanical and array scanning, where the TRXs in the array simultaneously scan a full row of the spot grid while a mechanical movement of the reflector provides coverage for the vertical dimension.

Within a mechanical scanning approach, different patterns to scan the FoV can be designed. Most scan patterns can be classified as a variation of three approaches: vertical raster scanning, diagonal scanning, or elliptical/circular scanning. These three techniques are illustrated in Fig. 1.13. Raster scanning is a historically-common approach, but has severe speed and power consumption disadvantages due to the necessary accelerations and stops to switch directions during the scan. To partially solve this, each diagonal sweep in a diagonal scan is implemented with a highly-eccentric elliptical motion to avoid hard acceleration changes. The main disadvantages in this approach are the overlapping in upwards and downwards sweeps, as well as the tilted FoV coverage. Lastly, an elliptical or circular scan uses less eccentric ellipses to cover the FoV vertically instead of horizontally, allowing for smoother, faster, and less consuming scanning. The main disadvantage of this approach would be the considerable overlap among ellipses.

For any of the scanning techniques a tradeoff must be made between the integration time per pixel (T_{pixel}), the number of pixels per image (N_{pixels}), and the acquisition time per image (T_{image}). These variables are related by the following equation:

$$T_{image} = N_{pixels} \cdot T_{pixel} \quad (1.25)$$

An ideal scan case would have very low T_{image} to allow fast image refresh, high N_{pixels} to provide good pixel density in the image (leading to better quality), and high T_{pixel} to achieve maximum SNR. This ideal scenario is limited by several technological constraints. For instance, T_{image} is lower-bound by the scanning approach, which in the case of mechanical scanning will depend on motor speed, and T_{pixel} is lower-bound by signal generation, which might not be able to produce accurate frequency sweeps at the desired rate.

Another architectural classification can be made regarding the relative placement of the target, TX, and RX. If the target is located between the transmitting and receiving antennas, the measured data correspond to transmissions through the target. If the receiving antenna is placed at same side of the target as the TX, then the measured data come from

Figure 1.13: (a) Vertical raster scanning. (b) Diagonal scanning. (c) Elliptical scanning.

reflected waves. Each type of data is particularly useful for different scientific interests: Transmission data can provide information about a target's dielectric constants and volumetric size, while reflection data give information about the target's surface and layers.

The two radars discussed throughout this thesis are both FMCW systems. For the 300 GHz radar in Chapter 2 the high-frequency architecture is implemented using COTS waveguide components, while for the 122 GHz system in Chapter 3 the front-end is a commercial MMIC. Both systems use a linear frequency sweep with approximately the same chirp duration (although different bandwidths), and similar baseband processing. Although both imagers use mechanical scanning, for the 300 GHz the scan pattern is elliptical, while for the 122 GHz system a diagonal scan is used. In addition, the 122 GHz imager presents the ability to use multi-TRX scanning in combination with its mechanical sweeps.

1.2.3 State-of-the-art Imaging Systems

This section is devoted to compiling some of the recent state-of-the-art security imagers in the mm-wave and submm-wave bands. These systems are summarized in Table 1.1.

Table 1.1: Recent state-of-the-art active and passive imagers for security screening

Reference	Type	Frequency (GHz)	Bandwidth (GHz)	Resolution *	Chirp Duration (μ s)	Commercial [†]	Real-Time [‡]
Our work [104]	Active	300	27	1.6x1.6x1.1	150	No	Yes
Our work [105]	Active	122	15	2.4x2.8x1.8	150	No	No
Yurduseven <i>et al.</i> [106]	Active	340	29	0.76x0.77x0.52	82.5	No	Yes
Robertson <i>et al.</i> [107]	Active	340	30	1x1x1	40.96	No	Yes
Sheen <i>et al.</i> [108]	Active	350	19.2	1x1xN/A*	N/A*	No	No
Ahmed <i>et al.</i> [109]	Active	75	10	0.2x0.2xN/A*	N/A*	Yes	Yes
Bryllert <i>et al.</i> [110]	Active	220	24	1x1x1	80	No	Yes
Gui <i>et al.</i> [111]	Active	340	28.8	N/A*xN/A*x0.7	300	No	N/A*
Stanko <i>et al.</i> [112]	Active	300	40	0.13x0.13x0.5	62	No	No
Cheng <i>et al.</i> [113]	Active	340	16	1.4x1.2x1	75	No	Yes
Grzyb <i>et al.</i> [114]	Active	240	60	0.1x0.1x0.25	2000	No	No
Mostajeran <i>et al.</i> [115]	Active	220	62.4	0.2x0.2x0.27	50	No	No
L3 ProVision 2 [116]	Active	27	5.75	N/A*	5.59	Yes	No
Smiths Detection ego [117]	Active	24	0	0.4x0.4xN/A*	0	Yes	Yes
ThruVision TS4-SC [118]	Passive	250	40	5x5xN/A*	N/A*	Yes	Yes
Asqella ARGON [119]	Passive	250 & 420 & 625	N/A*	2.5x2.5 & 0.4	N/A*	Yes	Yes
Heinz <i>et al.</i> [120]	Passive	350	80	2x2 & 0.4	N/A*	No	Yes

*Volumetric resolution (cm^3) for active imagers, and spatial (cm^2) and temperature (K) resolution for passive imagers.[†]Distinguish between commercial products and academic prototypes.[‡]Considered real-time if image acquisition and processing are faster than or equal to 1 fps.

*Not available or mentioned in referenced work.

1.3 Thesis Motivation and Objectives

This thesis details the improvement of a pre-existing 300 GHz standoff imaging radar and the development of a new hybrid imager with active and passive sensors, at 122 GHz and 94 GHz, respectively. The development of the new imager is a natural transition from the 300 GHz system, as RF technology progressively matures and allows integrated sensors to appear commercially in the mm-wave band. For both radars, the contributions performed involve all the fields necessary in modern mm-wave imaging systems, such as signal and image processing, RF component characterization, or mechanical control. Although the intended main application of both systems is security screening of people, several other applications, like data acquisition or vital sign monitoring, are also described and tested.

The objective of this thesis is twofold. First, improvements on several areas of the 300 GHz imager are desired in order to enhance its performance as far as image refresh speed, threat detection, and area scanned. To do so, modifications to the original hardware and software architectures must be designed, implemented, and evaluated. The second goal is to develop an imager able to simultaneously scan a scene with active and passive methods, thereby enhancing its detection capabilities in comparison to single-sensor systems. From an engineering point of view, a particular focus is placed on the developments to implement low-cost solutions, using COTS components whenever possible, in order to offer an economical alternative to traditionally expensive imaging systems in the mm-wave bands.

The work performed for this thesis contributes to technological development in the areas of mm-wave imaging and radar signal processing through the items summarized below:

1. Improvements to a pre-existing 300 GHz standoff imaging radar with the aim of enhancing its performance in several key areas.
 - Implementation of a real-time signal and image processing architecture, based on parallel computing in a Graphics Processing Unit (GPU) that allows refresh faster than 9 fps (detailing optimizations for additional speed).
 - Assessment of a multistatic configuration for the radar's TX and RX, through an image-based study of the scattering involved and field measurements of several scenarios.
 - Initial development and testing of an NDT-related data acquisition application, based on low-cost plastic guides, as well as a vital sign monitoring application to detect and track a human's respiration and heartrate.
2. Development of a hybrid imaging radar with 122 GHz active sensors and 94 GHz passive sensors.
 - Improvements to the original 122 GHz commercial radar in its RF and processing architectures.
 - Characterization of both sensors in standalone configuration and in simultaneous scanning.
 - Development of a processing algorithm for the images captured by the sensors based on image registration and fusion with an optical device.
 - Other minor developments, such as the testing of the additional applications initially developed for the 300 GHz radar.
3. Design and evaluation of automatic threat detection algorithms based on neural networks, for both active and passive images in simulation of security checkpoints.

1.4 Publications and Author's Contributions

Most of the results presented in this thesis have also been published in the following journals, conferences, and workshops (where the contributions of the thesis's author are highlighted):

1. C. Hernández, **F. García-Rial**, D. Montesano, J. I. Godino-Llorente, and J. Grajal, "Cardiopulmonary activity monitoring using a millimeter wave radar", *IEEE J. Electromagn., RF, Microw. Med. Biol.*, (pending submission).
 - The author was in charge of the radar's development and set-up's design, performed some initial testing and measurements, and wrote part of the publication's text.
2. **F. García-Rial**, D. Montesano, L. Perez-Eijo, M. Arias, B. Gonzalez-Valdes, A. Pino, and J. Grajal, "Evaluation of Standoff Multistatic 3-D Radar Imaging at 300 GHz", *IEEE Trans. THz Sci. Technol.*, (submitted for publication).
 - The author was in charge of the set-up's design, performed all measurements, and wrote the publication's text.
3. L. Perez-Eijo, B. Gonzalez-Valdes, M. Arias, D. Tilves, Y. Rodriguez-Vaqueiro, O. Rubinos, A. Pino, **F. García-Rial**, and J. Grajal, "A Physical Optics Simulator for Multireflector THz Imaging Systems", *IEEE Trans. THz Sci. Technol.*, (submitted for publication).
 - The author was in charge of performing the verification measurements and establishing their set-up.
4. **F. García-Rial**, D. Montesano, I. Gómez, C. Callejero, F. Bazus, and J. Grajal, "Combining Commercially Available Active and Passive Sensors into a mm-Wave Imager for Concealed Weapon Detection", *IEEE Trans. Microw. Theory Tech.*, vol. 67, no. 3, pp. 1167 - 1183, Mar. 2019.
 - The author was in charge of the active sensor's development, performed all measurements, and wrote the publication's text.
5. **F. García-Rial**, M. Ramírez-Torres, L. Úbeda-Medina, and J. Grajal, "FMCW Radar for Data Acquisition on High-Density Polyethylene Guides", *IEEE Sens. Lett.*, vol. 2, no. 3, Sept. 2018.
 - The author performed all measurements, participated in the design of the guides and set-up, and wrote the publication's text.
6. **F. García-Rial**, L. Úbeda-Medina, and J. Grajal, "Real-time GPU-based image processing for a 3-D THz radar", *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 10, pp. 2953 - 2964, Oct. 2017.
 - The author designed and implemented the processing architecture, performed all measurements, and wrote the publication's text.
7. J. Grajal, G. Rubio-Cidre, **F. García-Rial**, G. Ortiz-Jiménez, L. Úbeda-Medina, "Terahertz Radar Developments at UPM", *Workshop on Submillimetre Wave and THz Imaging for Security at the 2017 European Microwave Week*, Oct. 2017.

- The author performed some measurements, and wrote part of the publication's text.
8. G. Ortiz-Jiménez, **F. García Rial**, L. Úbeda Medina, R. Pagés Scasso, N. García Santos, and J. Grajal, "Simulation Framework for a 3-D High-Resolution Imaging Radar at 300 GHz with a Scattering Model Based on Rendering Techniques." *IEEE Trans. THz Sci. Technol.*, vol. 7, no. 4, pp. 1 - 4, July 2017.
 - The author was involved in the framework's design, performed some measurements, and wrote part of the publication's text.
 9. J. Grajal, G. Rubio-Cidre, A. Badolato, L. Úbeda-Medina, **F. García-Rial**, and A. Garcia-Pino, "Compact Radar Front-End for an Imaging Radar at 300 GHz, *IEEE Trans. on THz Sci. Technol.*, May 2017.
 - The author performed some measurements, and wrote part of the publication's text.
 10. **F. García-Rial**, L. Úbeda-Medina, J. Grajal, and O. Rubiños, "Optimization of a compact THz imaging radar for real-time operation", *2016 National Symposium of the URSI*, pp. 1 - 4, September 2016.
 - The author performed some measurements, and wrote the publication's text.
 11. **F. García-Rial**, L. Úbeda-Medina, J. Grajal, A. Badolato, G. Rubio-Cidre, A. Garcia-Pino, B. Gonzalez-Valdes, and O. Rubiños, "Optimization of a compact THz imaging radar for real-time operation", *2016 Global Symposium on Millimeter Waves & ESA Workshop on Millimetre-Wave Technology and Applications*, pp. 1 - 4, June 2016.
 - The author performed some measurements, and wrote the publication's text.
 12. J. Grajal, G. Rubio-Cidre, A. Badolato, L. Úbeda-Medina, **F. García-Rial**, G. Ortiz-Jiménez, "Terahertz Radar Technologies", *Spain-Japan Workshop on Millimeter-wave and Terahertz*, March 2016.
 - The author performed some measurements.

1.5 Research Projects Involved

The work performed during this thesis has been partially or fully funded by the following national research projects or grants:

- **Ministerio de Ciencia, Innovación y Universidades:** "Sensores heterodinos activos en la banda de terahercios para detección remota", TEC2011-28683-C02-01, 2012-2015.
- **Ministerio de Ciencia, Innovación y Universidades:** "Array digital en recepción para radares FMCW de vigilancia superficial", TEC2014-53815-R, 2015-2017.
- **Ministerio de Ciencia, Innovación y Universidades:** "Comunicaciones con transceptores en MIMO LTE-A/MMW para servicios avanzados en transporte de alta movilidad", TEC2017-87061-C3-1-R, 2018-2021.
- **Comunidad de Madrid:** "Radar para detección de basura espacial (SPADERadar-CM)", P2013/ICE 3000, 2014-2016

- **Ministerio de Educación, Cultura y Deporte:** “*Diseño y fabricación de radares en la banda de THz*”, Programa de Ayudas para la Formación de Profesorado Universitario 2015, FPU15/03095, 2016-2020
- **Dirección General de Tráfico:** “*AUTODRIVE: Diseño y construcción de un sensor de posicionamiento de alta disponibilidad, fiable y con precisión milimétrica para un sistema de conducción autónoma*”, SPIP2015-01879, 2015-2017

1.6 Thesis Structure

Aside from the initial introduction in Chapter 1, the thesis is divided into four additional chapters, with the two first parts separated according to the imagers involved. The organization is as follows:

- Chapter 2 deals with the developments made for the radar at 300 GHz, which is an imaging system composed of a single sensor type. The initial imager is first described in order to establish a base point for the following improvements. The next sections address the description and testing of the new real-time processing architecture and the multistatic configuration. Future improvements available in all areas are also detailed. Part of the work in this chapter is presented in publications 2 - 3 and 6 - 12, from Section 1.4.
- Chapter 3 is focused on the implementation of a hybrid system at 122 GHz and 94 GHz. In first place, the characterization of the performance of a 122 GHz radar is presented. The hybrid imager is then described and tested, with individual and simultaneous scanning from the sensors. Part of the work in this chapter is presented in publication 4, from Section 1.4
- Chapter 4 analyzes the introduction of a processing architecture for the automatic detection of concealed threats in active and passive images. This architecture is based on neural networks trained on public and in-house datasets of images simulating passengers in a security checkpoint. The design of the machine learning models and their performance are fully detailed, along with several approaches for improvement.
- Chapter 5 addresses some of the different applications for the two imaging systems, aside from security screening. Applications such as data acquisition from plastic guides and vital sign monitoring are tested with both radars. Part of the work in this chapter is presented in publications 1 and 5, from Section 1.4.
- Chapter 6 simply summarizes the conclusions derived from the work performed, and some of the future research lines that could be pursued for both systems.
- The Annexes complement the information included in the previous chapters. In these pages the characterization and datasheets of RF components from the 300 GHz imaging radar are included, as well as examples of the image formation options available and minor updates done outside of this thesis’ work. A detailed summary of the electronics and control stages used for synchronization during the radar’s real-time processing is also presented. The Annexes also include a summary of modifications made to a COTS 122 GHz radar kit, and a proposal for bandwidth enhancement for the 300 GHz radar. Additionally, documentation regarding the materials and equipment used for the data acquisition application is also attached. Radar analyses of several materials encountered during the thesis are also included. Lastly, proposed future work for the 300 GHz imaging radar, regarding the automatic recognition of threats and image quality, is described.

Chapter 2

300 GHz Single-sensor Security Imager

This chapter addresses the developments to a pre-existing 300 GHz FMCW imaging radar for standoff security screening. The system was designed and built during a previous thesis, but has since then received several improvements. In this chapter, some of these improvements are detailed. The chapter is organized into the following sections:

- Section 2.1 starts with an introduction to the 300 GHz radar, to provide some base knowledge of the system before describing the improvements made. The section details the radar's architecture, as well as considerations regarding the scanning mechanism and the imaging variations.
- The first improvement to the radar is described in Section 2.2, which presents the development of a real-time signal and image processing architecture based on GPU platforms. Two different iterations of this parallelized architecture are described, tested, and compared with the previous serial software.
- The second improvement is presented in Section 2.3, which describes a multistatic imaging configuration. This architecture is based on the use of an isolated TX to globally illuminate the scene and provide a bistatic propagation path to the 300 GHz radar. Through several placements of the TX, multiple bistatic and monostatic images can be combined to form a single multistatic result. An initial implementation of this configuration is designed and tested with multiple imaging scenarios.
- Lastly, Section 2.4 lists some of the future improvements and research lines available in relation to the 300 GHz radar.

2.1 Introduction to a 300 GHz Standoff Imaging Radar

The contributions presented in this chapter are exclusively focused on improving the performance of a single 300 GHz FMCW standoff 3-D imaging radar for security screening, developed at the *Grupo de Microondas y Radar* (GMR) of the *Universidad Politécnica de Madrid* (UPM). The development of this system was completed a year previous to the start of this thesis [102,121], but since then, the imager has continued progressing through several optimizations such as improved signal generation [97], real-time processing [104,122], reduced cost and size via a self-mixing TRX [123], or THz scattering analysis [22]. Some of the radar's

parameters most used throughout this thesis are summarized in Table 2.1. In this section, a brief description of the radar system is presented, along with an explanation of the signal and image processing involved. Additionally, a summary of some of the minor modifications performed to the system (not completely detailed in previous publications) outside of the thesis' objectives is provided in Annex A.4.

Table 2.1: Relevant system parameters for the 300 GHz imaging radar

Parameter	Value
Standoff distance (m)	8
Central frequency (GHz)	300
Transmitted bandwidth (GHz)	27
Transmitted power (dBm)	~ -1
Typical chirp period (μ s)	150 [‡]
Field of View (cm ²)	50x90
Spot size* (cm)	1.6
Range resolution [†] (cm)	1.1
Scan time (s)	≥ 0.55
Sampling frequency (MHz)	40
Power consumption* (W)	16.6

[‡]Usual value, but can be lower.

*At -6 dB and approximately symmetric in X and Y dimensions.

[†]After windowing.

*For mm-wave components only.

2.1.1 System Description

The imager is implemented on a homodyne architecture, chosen for its reduction in cost and size compared to a heterodyne architecture¹, based on frequency upconversion from a microwave source. A DDS is used as the chirp source², generating a signal with 500 MHz bandwidth between 500 MHz - 1 GHz. The chirp is then upconverted using a 4.8 GHz signal generated from a Phase-Locked Oscillator (PLO) locked to a 10 MHz master oscillator which also controls the DDS and the sampling of the received signal, thereby ensuring coherence. Afterwards, the chirp is conditioned (filtered, amplified, and tripled) and directed to an SMA output between 15.9 GHz - 17.4 GHz. From this connection, a power divider is used to feed the two RF chains available: one for the TX, and one for the RX.

On the TX side, an active x6 multiplier is used to shift the chirp to 95.4 GHz - 104.4 GHz. This component is followed by a medium-power amplifier (MPA), with an isolator before it to improve adaptation between the two devices. The final multiplication stage is performed by a x3 multiplier, generating the transmitted bandwidth of 286.2 GHz - 313.2 GHz. A 300 GHz central frequency has been chosen as a compromise between the better transmissivity through clothing of lower frequencies, and the increased cross-range resolution that comes with smaller wavelengths [20]. The TX ends with a radiating element in the form of a Picket-Potter horn antenna with a 12° Half-Power Beamwidth (HPBW) and 23 dB gain, easier to machine than corrugated horn antennas and with similar performance.

On the RX side, the chirp first reaches a x3 multiplier, moving the signal to 47.7 GHz - 52.2 GHz. Afterwards, another MPA is used to amplify the chirp, before reaching a manual attenuator necessary to progressively increase signal power to the following components during system start-up. After the attenuator another x3 shifts the chirp to 143.1 GHz - 156.6

¹Where additional Local Oscillators (LO) are needed.

²Due to its excellent frequency resolution, sweep speed, and linearity [124].

GHz, and directs it through an isolator. The last $\times 2$ multiplication factor for mixing with the received signal is provided by a subharmonic mixer (SHM). The signal is received through another Picket-Potter horn antenna identical to the TX side. The antennas are separated by a silicon beamsplitter to improve isolation, due to the lack of wideband circulators at the desired band. Only 56% of the incident power passes through the beamsplitter (with the other 44% being reflected with a 90° offset), in both transmission to and reflection from the target. A photograph of the two RF chains is presented in Fig. 2.1. It is worth noting that the micro- and mm-wave components used are all intentionally COTS, in an effort to provide a lower-cost alternative to traditional custom-made imagers.

Figure 2.1: Photograph of the TX and RX mm-wave chains of the 300 GHz imaging radar.

At the IF port of the SHM, the beat frequencies from mixing the received signal are generated. The IF signal is amplified and bandpass filtered (to eliminate leakage but also high-frequency noise) before reaching an ADC for sampling. The ADC is integrated in a conventional PC, which controls DDS signal generation, ADC sampling, and the radar's scanning (this last step through custom-made C++ code). The sampled data are then processed with MATLAB software to generate 3-D images from each scan.

To focus radiation at the desired 8 m range, the radar uses a Bifocal Ellipsoidal Gregorian Reflector System (BEGRS) designed by the *Universidad de Vigo* [125]. The main components of this system are: (a) a 60-cm main reflector to provide gain, (b) an oversized subreflector to scan a large FoV, (c) a moving mirror which creates the scanning pattern, (d) a feed reflector to illuminate the subreflector with a tapered collimated beam, and the aforementioned (e) beamsplitter and (f,g) horn antennas. The reflectors, as well as the table structure in which they are placed, have been machined from aluminum alloy 6082. With the use of the reflector system, targets 8 m away from the radar will produce a maximum in received power, while any target in a different range will be unfocused. This effect can be seen in Fig. 2.2, where the target is a metallic sphere, which registers a focusing range near 10 m (instead of 8 m) due to the propagation path of the signal through the reflector system.

The scanning itself is performed via a two-axis mechanical movement of the mirror, through the use of two motors: a rotation motor that spins the mirror in a fast movement, and an elevation motor that tilts it in a slow movement. The mirror's rotation causes the focusing spot at 8 m to depict ellipses (~ 50 cm wide and ~ 40 cm high), which are then vertically shifted with the tilting motion, effectively scanning a 50×90 cm² FoV centered at a height of 1.147 m from the floor with a ~ 1.6 cm circular spot. This elliptical trajectory was chosen instead of other conventional raster scan methods to avoid sudden accelerations, which degrade the uniformity of the scan [126]. With this scanning mechanism the FoV is swept at a minimum time of 0.55 s. A photograph of the reflector system and RF chains, as

Figure 2.2: Spectral comparison of a focused and unfocused metallic sphere with the radar using the BEGRS reflector system.

well as diagrams of the mirror movement and scanning pattern created, are shown in Fig. 2.3.

A block diagram of the radar, divided into four subsystems (signal generation and sampling, microwave, mm-wave, and reflectors), is shown in Fig. 2.4.

In addition to the introduction provided, Annex A.1 describes the different image representation options in the radar, Annexes A.2 and A.3 propose improvements regarding automatic recognition of targets and image quality, respectively, and Annex A.4 details minor system updates not included in this thesis.

Although the 300 GHz imaging radar is already at a state-of-the-art level in regard to qualitative image quality and refresh speed, some improvements can still be made in other system areas. Newer standoff imaging radars in the same frequency band, have demonstrated their ability to transmit bandwidths larger than the 27 GHz currently used by our system. This increase in transmitted bandwidth is of course beneficial to image quality and threat detection by producing finer range resolution. With the goal of enhancing the transmitted frequency range in the 300 GHz imager, several alternatives are proposed in Annex G, including solutions where no updates or substitutions to the RF equipment are necessary. Furthermore, as is detailed in Chapter 3, modern RF systems are trending towards more compact and less expensive devices with the use of MMICs and other integrated technologies. The application of these techniques to reduce an imaging radar's size and cost would be advantageous for both its industrial fabrication (lowering material and handling costs) and its on-field performance (improving mobility and concealment). For this purpose, also in Annex G, a proposal for size-reduction in the 300 GHz radar's RX mm-wave chain, suitable for its compact-imaging configuration, based on COTS MMIC products is described [127].

Figure 2.3: Photograph detailing the reflector subsystem, diagram of the mirror's movement, and representation of the scan pattern over a human target, with the axes used throughout the text. The z axis goes through the torso, from front to back.

Figure 2.4: Block diagram of the 300 GHz imaging radar, divided into its main subsystems.

2.2 GPU-based Real-time Image Processing

Threat-detecting active imaging radars have proven to be apt in revealing hidden person-borne threats, by way of the images produced in each scene scan [50, 102, 128–131]. In the initial development stage, the main focus of these radar systems is not often placed at the processing time required for each scan. Nevertheless, this benchmark is of utmost importance in real-life implementations of security systems. The time required to scan and process the scene poses a limit to how fast a threat can be detected and removed. In an ideal scenario, the operator would be presented with a video sequence at a high enough frame rate to adequately examine moving targets [126].

Adjusting to these real-time requirements has lately been one of the top priorities of standoff imaging radars. For example, NASA’s Jet Propulsion Laboratory (JPL) has reported increasing performance improvements in the past years [128, 132] due to the use of a multi-TRX architecture and a Field-Programmable Gate Array (FPGA) to perform signal processing. Researchers from the Millimetre Wave Laboratory of Finland have also proposed the use of a multi-TRX scheme, by means of a patch-antenna array with electronic beam steering enabled by a FPGA [133]. F. Friederich et al. [134] coincide in choosing an array architecture, although this time with a GPU as a general-purpose processor, obtaining 40x10 cm² images in 200 ms.

A comparison among some state-of-the-art security screening radars is presented in Table 2.2, with the GMR’s radar including the work done in this thesis. The pixels (detections) to be processed per second can give a measure of the parallelization and speed achieved in each system. It is also necessary to take into account the way these pixels are displayed. For instance, filtering them or reconstructing surfaces requires more processing than displaying the unmodified point cloud.

In its initial stage, the speed of the GMR’s 300 GHz imaging radar was limited by two main factors:

- **Hardware limitations:** The radar’s scanning mechanism (shown on Fig. 2.3) presents a minimum image acquisition time of 0.55 s, which would allow a nearly 2 fps image refresh, if all data processing was done in parallel. The limitation comes from the maximum speed of the motors used to move the mirror. A more detailed review of the motor limitations is presented in Annex A, although mechanical improvements have not been performed during this thesis. Slightly higher image acquisition times are possible with newer versions of these motors, but a more impactful improvement could be achieved by introducing multi-pixel scanning [135], or avoiding mechanical scanning altogether [138, 139].
- **Software limitations:** The software used to process the radar’s acquired data was implemented on a MATLAB environment, and was executed on an offline manner. Due to these conditions, processing and representation of the radar data required more than 2 minutes per image. The offline nature of the code restricted the possibility of a video sequence, while the required processing time per image limited the refresh rate of said sequence. These limitations have been surpassed with the work here presented.

The software restrictions of the radar are due to the slower processing times of MATLAB versus other programming languages, and of serial execution versus parallelization schemes. Both of these obstacles can be overcome by migrating the signal and image processing architecture towards a parallel programming model in a different language. Recently, the plateauing of commercial CPU clock speeds has motivated a shift in computing applications from traditional serialized architectures to parallel multi-core schemes. Multiple hardware

Table 2.2: Comparison of real-time active imaging radars for security applications

Reference	FoV (cm ²)	Frames/second (fps)	Pixels/scan	Pixels/second (pps)	Hardware [†]	Image display [⊙]
Our work	50x90	9.42*	6000	56520	GPU	Reconstruction
Cooper [135]	15x21	8.3	2496	20716.8	2 FPGAs	Filtered point cloud
Robertson [136]	50x50	10	1296	12960	GPU	Point cloud
Robertson [92]	100x100	7	14400	100800	(CPU only)	Interpolated point cloud
Bryllert [110]	70x70	1	12500*	12500	FPGA	Point cloud
Cheng [137]	60x60	5	1600	8000	GPU + FPGA	Contour plot
Cheng [113]	60x200	1.43	12000 [‡]	17160	4 FPGAs	Interpolated point cloud
Ahmed [109]	200x100	0.5	3072	1536	Industrial PC + 8 FPGAs	Reconstruction

[†]Not including a conventional PC or CPU.

[⊙]Point clouds are the base representation. Filtered clouds and contour plots add a 2-D or 3-D filter to the representation, while interpolated clouds create a grid with interpolated data among captured coordinates. Reconstructions generate a surface from the spots in a cloud.

*Overlooking mechanical restrictions (which limit refresh to 2 fps).

* Estimated from chirp duration and frame rate.

[‡] Estimated from total angle coverage and motor steps.

alternatives have arisen to adapt existing systems in this endeavour. Some of the more common ones are multi-core processors, FPGAs [140], Digital Signal Processors (DSP) [141], and GPUs [142].

For the 300 GHz radar's new parallelized algorithm, a hybrid CPU-GPU solution has been chosen, where data are offloaded from the CPU to the GPU to be processed, in a paradigm called General Purpose computing on GPU (GPGPU). The low cost, ease of use, and optimization for graphics applications of these devices have been the primary reasons for this decision.

Some serial instructions in the algorithm remain in the CPU, while the parallelized functions have been implemented on NVIDIA's Compute Unified Device Architecture (CUDA) platform using C++ code [143]. CUDA is a paradigm developed to program NVIDIA's GPUs in a simple fashion, concealing some of the more intricate notions of parallel programming. It is based on a Single Instruction Multiple Thread (SIMT) model, in which groups of threads executed on the GPU's cores can process the same instruction on different data. Threads are grouped into two levels of 2-D matrices, where the first level (called *grid*) contains a second level (called *block*) in each of its cells. CUDA functions (called *kernels*) execute a grid of threads (scheduling them in groups of 32, called *warps*), which can access the different types of memory available to the GPU (per-thread local, per-block shared, and global). The organization of CUDA threads and their memory in a GPU may be seen in Fig. 2.5, while several descriptions of the CUDA framework may be found in the literature [144–146].

Figure 2.5: CUDA thread (*left*) and memory (*right*) organization. Reproduced from [143]

Two GPU models have been used during this thesis, both from NVIDIA: a GeForce GTX 970 (used in Section 2.2.1), and a TITAN X (used in Section 2.2.2). Their specifications may be seen in Table 2.3.

Some of the radar's previous modifications [123] allowed reduced volume, cost, and power consumption, paving the way towards a future industrial prototype, but they did not enhance the scanning mechanism or signal processing. In this section the design and operation of a new hybrid GPU-CPU-based algorithm which has made 9.42 fps image processing possible in the radar system will be detailed. First, the initial signal and image processing architecture is presented in Section 2.2.1. After detailing some of the initial design's processing limitations, the final architecture is presented in Section 2.2.2. Sample sequences of real-time imaging scenarios captured with the radar are shown in Section 2.2.3. Some of the ongoing and

Table 2.3: GPU Specifications

Parameter	NVIDIA GTX 970	NVIDIA TITAN X
CUDA Architecture	Maxwell	Pascal
CUDA Compute Capability	5.2	6.0
Clock Speed (MHz)	1178	1417
Global Memory (GB)	4	12
Memory Bandwidth (Gbps)	1792	3360
Streaming Multiprocessors	13	30
CUDA Cores	1664	3840
Shared Memory (KB)*	96	3000
Warp Schedulers*	4	2
ALU Units*	128	128
Special Function Units*	32	N/A

*Per Streaming Multiprocessor.

future work to enhance the processing scheme is later described at the end of this chapter, in Section 2.4.

2.2.1 Initial Processing Architecture

Since the ideal IF signal in a FMCW radar is a sinusoidal function of f_b frequency, the most common approach to obtain the beat frequency is using a discrete FFT [147] on the signal's samples. This algorithm allows a frequency-domain representation of the IF signal, and is the technique implemented on the GMR's screening radar. The FFT can be seen as a filter bank that allows separation of a scene's range (R) information in bins.

The signal and image processing algorithm initially implemented followed the sequence shown in Fig. 2.6, and detailed in Fig. 2.7.

Figure 2.6: Main stages of the initial signal processing algorithm, in execution order from left to right, and top to bottom.

The main innovation from previous versions of the processing architecture comes from the signal conditioning: to speed up the spectrum computation and peak detection in each ramp, a downsampling stage is performed. This step reduces the number of samples in each chirp period, thereby lowering the FFT length and processing time. The processing stages, each implemented in its own kernel, can be described as follows:

- **Calibration.** The transmitted signal is deteriorated by nonlinearities in the micro- and

Figure 2.7: Scan pattern of the GMR's imaging radar and sample results from some of the signal processing events in the algorithm. Notice that the x -axis in (c) has already been labeled according to R instead of frequency, through the relationship seen in (1.18). The final product is a 3-D point cloud of the internal layer of a female mannequin wearing a T-shirt and hiding a simulated explosive device in the abdomen area. Colorbar: Range in meters (R -axis).

mm-wave components³. A procedure to calibrate the received signal ($s[n]$, of length N) taking into account these effects [126] consists of correcting the signal with a reference received waveform from a point-like target ($r[n]$, of length N). A window function is also included in the reference waveform, in order to maximize dynamic range [95]. Both techniques can be summarized as a series of multiplications on the samples of the received signal, producing a calibrated signal ($s_c[n]$) such as

$$s_c[n] = s[n] \cdot r[n], \forall n \mid 1 \leq n \leq N, s[n] \in \mathbb{R} \wedge r[n], s_c[n] \in \mathbb{C} \quad (2.1)$$

- **Frequency Shift.** A baseband shift is necessary before downsampling the RX signal. The kernel is completely analogous to the calibration/windowing kernel, with the difference being the reference signal used in the multiplication. In this case, the signal is a complex exponential with the appropriate frequency to do the shifting. Let f_{shift} be the beat frequency associated to the standoff distance desired, then the shifted signal ($s_s[n]$) is calculated as

$$s_s[n] = s_c[n] \cdot e^{-j2\pi n f_{shift}}, \forall n \mid 1 \leq n \leq N, s_s[n] \in \mathbb{C} \quad (2.2)$$

- **Low-pass Filtering.** Filtering is necessary to avoid aliasing when downsampling. This kernel has similar operations to the two before, with a different reference signal⁴. The filter has been previously designed to allow passing of the beat frequencies corresponding to 30 cm around the standoff distance. The filtering is performed in the frequency domain. Let S_k and H_k be the N -length Fourier transforms of $s_s[n]$ and the low-pass filter, respectively. Then the filtered signal ($s_f[n]$) can be calculated as

$$s_f[n] = \mathcal{F}^{-1}(S_k \cdot H_k), \forall n, k \mid 1 \leq n, k \leq N, s_f[n] \in \mathbb{C} \quad (2.3)$$

- **Downsampling.** This stage is tasked with reducing the number of samples in the chirp in order to speed-up further processing. In the time domain, this can simply be done

³Such as non-constant efficiencies and gains, temperature-related variations, and impedance mismatches.

⁴In this case the reference is a pre-designed filter read from a file.

by selecting the desired samples out of the total N , and discarding the rest. Assuming an integer decimation factor M , where N is divisible by M , the filtered signal ($s_f[n]$) becomes

$$s_d[n] = s_f[Mk], \forall n, k \mid 1 \leq n \leq \frac{N}{M} \wedge 1 \leq k \leq N, s_d[n] \in \mathbb{C} \quad (2.4)$$

- **Fourier Transform.** The main tool in this stage is the FFT, enabling detection of the frequency components of the received signal, as has been explained. Each frequency ramp will later correspond to an individual detection zone in the scanned area, so an FFT per chirp is necessary. Although intermediate FFTs can be included throughout the processing pipeline (to perform a specific operation in the frequency domain), the transform in this stage is the main and final FFT, through which the chirp's spectrum is extracted.
- **Spectrum Computation.** A representation of the RX signal spectrum in a logarithmic scale (S_{d_k}) is done, per frequency ramp, so that the range information (R) can be extracted. In order to speed up this stage, the spectrum can be calculated only for the R interval desired, where the target is located. This technique is called range gating (limiting, with maximum and minimum distances, the R axis), and is illustrated in Fig. 2.8. Since the ADC measures $s[n]$ in volts over a 50Ω impedance and the spectrum is to be represented in dBm, S_{d_k} is calculated by

$$S_{d_k} = 20\log(|\mathcal{F}(s_d[n])|) + 10\log\left(\frac{1000 \frac{mW}{W}}{2 * 50 \Omega}\right) \quad (2.5)$$

$$\forall n, k \mid 1 \leq n \leq \frac{N}{M} \wedge 1 \leq k \leq \frac{N}{2M}, S_{d_k}[n] \in \mathbb{R}$$

where only the positive-frequency side of the spectrum is considered.

Figure 2.8: Range gating (R -axis) in a target's silhouette.

- **Peak detection.** As can be seen in Fig. 2.7, the resulting spectrum is not an ideal sinc due to the presence of noise, and multiple targets and layers of clothing, thus containing more than one frequency peak. In order to isolate, for instance, a desired clothing layer or object, a peak detection algorithm is run on the frequency-domain signal. The implemented algorithm can distinguish between (a) the most powerful peak, (b) the external clothing layer, and (c) the internal (skin and hidden objects) layer. This algorithm may be run more than once for each image, in any combination of its three

settings, to display more than one range layer simultaneously. As an example, the peak detection algorithm to extract L range layers can be summarized by the diagram presented in Fig. 2.9.

Figure 2.9: Flow diagram of the peak detection algorithm for extracting range layers from a chirp's spectral samples.

To extract the outermost layers the second *for* loop in the algorithm must simply be iterated starting from the first spectral samples. The algorithms to detect the innermost and most-powerful layers are similar. For the former, the second *for* loop is simply iterated backwards (starting from the last spectrum samples), and for the latter, an additional loop is added to sort the detected peaks according to returned power.

- **Surface coordinates.** Each peak detection algorithm provides the range information of all chirp periods, which corresponds to the z coordinate on a 3-D Cartesian representation whose origin has been assigned to a fixed corner of the radar's mechanical structure. The x and y coordinates depend on the position of the mirror during the scan. These coordinates are obtained from custom polynomial functions [125] of the motors' elevation and rotation angles (interpolated from known starting and ending positions during a scan, assuming constant velocity). These functions are as follows

$$x = a_0 + a_1u + a_2u^2 + a_3v^3 + a_4u^3 + a_5uv^2 \quad (2.6)$$

$$y = b_1v + b_2uv + b_3u^2v + b_4u^3 \quad (2.7)$$

where $a_0, \dots, a_5$ and $b_1, \dots, b_4$ are constant coefficients obtained from least squares fitting, and u and v can be defined as

$$u = -\sin(\beta) \cdot \cos(\phi) \quad (2.8)$$

$$v = \sin(\gamma) \cdot \cos(\beta) + \cos(\gamma) \cdot \sin(\beta) \cdot \sin(\phi) \quad (2.9)$$

with γ and ϕ being the elevation and rotation angles respectively (as defined in Fig. 2.10), during a given moment in the scan, and β a constant 4.557° tilt of the mirror with respect to its surface normal.

Figure 2.10: Diagram of the elevation and rotation axes, as well as angles, in the radar's scanning mirror. Reproduced from [125].

The individual result of this stage is an (x, y, z) tuple per chirp period and per peak detection algorithm, which can be joined into a 3-D point cloud of detections, with the scan pattern seen in Fig. 2.7(a). There will be as many complete point clouds as peak detections run.

Since the most common detection processing approach produces a point cloud, the display of these results is usually centered around this representation [50, 126, 129, 130]. Additional image processing, such as interpolations, surface smoothing, and median filtering is often useful to compensate cavities and outliers in the point cloud.

In our system the user may choose to also execute a reconstruction algorithm on the 3-D point cloud, based on a Delaunay triangulation. The result is a reconstructed surface with triangular facets joining the detections of the original point cloud. Both of these representations (point cloud and reconstructed surface), which are the final results of the complete processing algorithm, are shown in Fig. 2.11.

Parallelization of the algorithm on a GPU

The signal processing algorithm described presents data levels⁵ in which processing is independent among groups of data, making it a good candidate for parallelization. In this regard, GPU threads have been tasked with the processing of individual samples in all kernels, except for peak detection. For the peak detection kernels, threads have been tasked with the processing for entire ramps. The resulting thread structure, for these last kernels, can be seen in Fig. 2.12, where the variables *gridDim* and *blockDim* give the GPU block and grid sizes respectively, in each axis.

Due to this configuration, individual linear frequency modulated (LFM) ramps or samples (C_i), depending on the kernel, can then be addressed in each kernel by thread indexing with the following equation

$$C_i = (bI.x + bI.y * gD.x) * (bD.x * bD.y) + (tI.y * bD.x) + tI.x \quad (2.10)$$

where *bI* stands for *blockIdx*, *tI* stands for *threadIdx*, *bD* for *blockDim* and *gD* for *gridDim*, with all four variables being native parameters of CUDA. The variables *bI* and *tI* represent the block position in a grid axis and the thread position in a block axis, respectively.

⁵Such as the complete frequency ramps or individual samples.

Figure 2.11: Sample point cloud and reconstructed images of the internal layer of a female mannequin wearing a T-shirt and hiding a handgun in the abdomen area. Colorbar: Range in meters (R -axis).

Figure 2.12: Relationship between data and thread structures in the parallel algorithm for a sample image. Each point cloud spot represents one LFM ramp, and is assigned to one GPU thread. Thread blocks are formed by grouping these spots in sub-images, as seen in the sample dotted squares.

The signal processing algorithm of Fig. 2.6 has been implemented in a GPU-based parallel scheme, while the display of the resulting image is still performed in CPU hardware (although now modified to execute in a parallel manner also). The entire sequence of events in an image scan and processing can be seen in Fig. 2.13.

Figure 2.13: Sequence of events and transfers, in chronological order from left to right, from the reception of the beat frequency until the display of its image. The legend details the connections in each stage, and the color code separates hardware platforms.

The synchronization of these events in a sample scan can be seen in Fig. 2.14, where an image is displayed by the CPU and then refreshed when the next image is finished processing. The developed architecture is ultimately limited by the mechanical scanning, which can only provide image data at a maximum speed of 0.55 s per image. Notice that GPU processing takes up much less time than 0.55 s, which means that refresh could be faster if the chirp samples were available sooner (i.e. if mechanical scanning was faster). GPU processing for a single image is divided into two steps due to limitations in memory transfers discussed in the following subsection.

Figure 2.14: Timeline of synchronization between scanning, image processing, and display in a sample scan. In this example, image refresh is limited by the mechanical scanning; processing can be contiguous and display can be faster, if the ADC data of all images are available sooner on the GPU.

2.2.1.1 GPU-code Profiling

The images considered in our development are made up of ~ 6000 chirp periods⁶, each with 6000 ADC samples (produced by using a $150 \mu\text{s}$ chirp and a 40 MS/s sampling frequency). As the radar scans the scene, the ADC is continuously sampling and buffering the RX signal. Data blocks from the ADC to the GPU have been arranged to hold 3000 LFM ramps each, resulting in two blocks per image. The effect of varying data size in blocks, for the developed algorithm, may be seen in Fig. 2.15, with configurations detailed in Table 2.4.

Figure 2.15: Average GPU-based signal processing rate per 6000-LFM ramp image for various ADC data sizes. Configurations vary depending on signal processing and display choices seen in Table 2.4. Each configuration time is averaged over 100 images, per datapoint in the x -axis. (a) Shows the required GPU processing time per image. (b) Shows the processing speed in images per second [the inverse of (a)].

Table 2.4: CUDA signal processing configurations

Config.	N ^o of peaks	Range layer chosen	Calibration	Time (s)*
1	1	External	Yes	0.2468
2	1	External	No (windowing only)	0.2477
3	1	Max. peaks	Yes	0.2476
4	1	Max. peaks	No (windowing only)	0.2485
5	1	Internal	Yes	0.2477
6	1	Internal	No (windowing only)	0.2476
7	2	Ext. and Max.	Yes	0.2498
8	2	Ext. and Max.	No (windowing only)	0.2494
9	2	Ext. and Int.	Yes	0.2489
10	2	Ext. and Int.	No (windowing only)	0.2496
11	2	Int. and Max.	Yes	0.2503
12	2	Int. and Max.	No (windowing only)	0.2503

*Per image, averaged over 100 images, and using 3000-chirp blocks.

The processing time in Fig. 2.15 encompasses all the GPU-based stages in Fig. 2.13. This, of course, includes the appropriate GPU kernels, but also memory transfers and CPU serial control code. It does not include image display code, which will be discussed later on. The different configurations considered comprise the signal processing variations shown in Table 2.4.

⁶Which means there are 6000 possible spots in the point cloud.

The average processing time with 3000-LFM-ramp data blocks for these configurations can be seen in the rightmost column. In the last six configurations, two range layers, instead of only one, were simultaneously displayed (Table 2.4). This requires executing an additional CUDA kernel, for peak detection of the second point cloud. Nonetheless, as can be seen from Table 2.4 and Fig. 2.15, it does not significantly influence the processing time required. This is due to the fact that the time required to execute an additional peak detection kernel is very reduced. It can be noticed that the remaining signal processing choices also do not deeply impact execution time. As can be seen in Fig. 2.15, for any given number of chirps per block the GPU-based algorithm performs better than MATLAB's 2-minutes-per-image mark. On the other hand, GPUs are optimized to handle large amounts of data, where the repercussion of parallelizing execution can truly be noticed. Thus, for a small number of LFM ramps in each packet, GPU occupancy is low and its performance is hindered. Furthermore, it may be noticed that the sharp decline in processing time from 1 to 500 LFM ramps per block is not present in the 500 to 3000 LFM ramps range. The asymptotic tendency of the curve indicates that further GPU occupancy, past a given number of chirps per block, does not improve performance.

The profiling results of all kernels involved in the processing architecture may be seen in Fig. 2.16 and Table 2.5. The profile was done with NVIDIA Visual Profiler [148], and the results are once again averaged over 100 images. FFTs are done in batched mode, which means multiple FFTs (associated with different ramps) are executed with a single instruction of CUDA's CUFFT library.

Figure 2.16: GPU processing time per image for all kernels, averaged over 100 images. Only one peak detection algorithm is chosen per image. FFTs are done in batched mode with CUDA's CUFFT library, while all other kernels are custom-made. Profiling is done with NVIDIA's Visual Profiler.

Notice that some of the kernels listed in Fig. 2.16 and Table 2.5 match the signal processing stages shown in Fig. 2.6, but there are also new functions. Added functions, such as the sinusoids kernel or the intermediate FFTs, are simply due to requirements from our specific implementation of the general processing architecture. All the kernels can be described, in execution order, as follows:

- **Sinusoids.** This kernel is tasked with computing and storing in GPU memory the sinusoids that make up the frequency-shifting exponential. The exponential function is variable according to the standoff distance selected by the user.

Table 2.5: Distribution of image time over kernels

Number	Kernel	Time (ms)	Time (%) [†]
1	FFT (before downsampling)	35.91	23.9
2	IFFT	39.98	25.4
3	Frequency shift	18.17	11.5
4	Spectrum computation	13.66	8.7
5	FFT (after downsampling)	23.61	13.9
6	Filtering	11.28	7.2
7	Calibration	8.48	5.4
8	Peak detection (external)*	4.29	2.7
9	Peak detection (internal)*	3.39	2.2
10	Peak detection (maximums)*	6.32	4.0
11	Downsampling	1.60	1.0
12	Polynomial	0.04	0.0
13	Sinusoids	0.00	0.0

[†]Percentage over sum of total GPU kernel time for an image.

*Only one of these kernels is chosen per image.

- **Calibration.** This kernel is tasked with the calibration of the IF signal, and its windowing. The mathematical operations are performed in the time domain. Data are stored in CUDA's *cufftDoubleComplex*, common for complex numbers, which jointly stores two *double* sized samples: one for the real part and one for the imaginary part. This data type is used throughout most kernels, even when data are real-only, in preparation for FFTs performed later on. This kernel corresponds to the calibration stage in Fig. 2.6.
- **Frequency shift.** This kernel is analogous to the calibration kernel, only varying the reference signal used for multiplication against the IF samples. The reference is, in this case, the exponential resulting from the sinusoids kernel. The multiplications are performed in the time domain. This kernel corresponds to the frequency shift stage in Fig. 2.6.
- **FFT (before downsampling).** In this kernel an intermediate FFT is performed (between the frequency shift and filtering stages in Fig. 2.6), different from the main Fourier Transform stage. It is responsible for transferring the IF signal to the frequency domain, before filtering. This is done to simplify the filtering kernel, which can then be reduced to multiplications instead of convolutions. The FFT is zero-padded to the next power of two of the frequency ramp size for faster performance⁷, since most FFT software libraries are optimized for power-of-two lengths.
- **Filtering.** This kernel is tasked with low-pass filtering the shifted and calibrated IF signal, in order to avoid aliasing before downsampling, and corresponds to the filtering stage in Fig. 2.6. The operations are performed in the frequency domain by multiplying against a reference filter stored in GPU memory.
- **IFFT.** This kernel performs another intermediate transform (between the filtering and downsampling stages in Fig. 2.6), in this case an Inverse FFT (IFFT) on the low-pass filtered IF signal, in order to shift back into the time domain⁸. This is done to simplify the downsampling stage, which then consists simply on re-indexing the signal samples.
- **Downsampling.** This kernel reduces the number of samples stored for each frequency ramp. Samples are dropped, according to the downsampling factor required, by skip-

⁷Since each chirp has 6000 samples, this means that the FFT has a length of 8192.

⁸The length of this IFFT is also 8192.

ping the associated array indexes when re-storing the frequency ramp. The downsampling factor (M) is related to the desired range window (R_{window}) through

$$M = \text{floor}\left(\frac{R_{max}}{R_{window}}\right) \quad (2.11)$$

where R_{max} is the maximum detectable range and the $\text{floor}()$ function is used to ensure that M is an integer⁹. This kernel corresponds to the downsampling stage in Fig. 2.6.

- **FFT (after downsampling).** This kernel is the main Fourier Transform stage in Fig. 2.6. It is the last FFT function executed¹⁰, and is tasked with generating the spectrum where the peak detection will later be performed.
- **Spectrum computation.** This kernel is tasked with modifying the result of the downsampled FFT to a logarithmic scale referenced over a 50 Ω resistance¹¹ in dBm, corresponding to traditional spectrum representations. For high-level optimization, the module operation and the complex sum [in the formula presented in 2.2.1], are performed with CUDA's *cuCabs* and *make_cuDoubleComplex* instructions. Since this computation takes the absolute value of the spectrum no complex numbers are necessary afterwards, and the resulting data are stored in C++'s *double* format. This kernel corresponds to the spectrum computation stage in Fig. 2.6.
- **Peak detection.** This kernel is responsible for selecting the R -axis layers that will be represented in the image. The kernel sorts the spectrum samples according to the peak criteria desired (i.e. to detect the maximum peak, the samples are ordered from most to least powerful). It corresponds to the peak detection stage in Fig. 2.6.
- **Polynomial.** This kernel is tasked with executing the polynomial functions that generate the (x, y) surface coordinates associated to each frequency ramp. It corresponds to the (x, y) calculation stage in Fig. 2.6.

The sum of kernels 1-7, 11-13, and one of 8-10, in Fig. 2.16 would be the total processing time required per average image. Performing this sum, taking the mean of kernels 8-10 instead of just selecting one, the total time per image is just 0.154 s. This total is approximately 0.10 s less than the results seen in Fig. 2.15. This is due to the fact that in Fig. 2.16 only GPU code has been profiled; the difference is due to CPU-executed control code, such as memory transfers. CUDA memory transfers (with CUDA's native function *MemCpy*) from CPU to GPU have been profiled at 0.08446 s (84.46% of the difference between Fig. 2.16 and Fig. 2.15). While pinned memory is already being utilized wherever possible, reducing the number and size of transfers and using asynchronous processing could still provide reduction of this overhead. In summary, considering only GPU-executed code (0.154 s per image), refresh could be as fast as 6.5 fps.

From these results we can establish some preliminary conclusions. First off, just the IFFT accounts for 25.40% of the total GPU time per image. Adding the FFT before downsampling and the one after downsampling, CUFFT functions now make up 63.20% of the total GPU time. If this percentage is distributed between FFT planning (instructions the GPU requires before performing the transform, according to CUFFT structure) and FFT execution, the results would be 22.80% and 40.40% respectively. The FFT after downsampling is, of course, much faster due to the reduced number of samples in the signal, which implies less operations. The direction of the FFT (direct or inverse) does not have such a considerable impact on

⁹With the usual scanning parameters, M has a value of 20.

¹⁰Its length is 4096. This value is the next power of two after applying M to 65536, which is the standard FFT length used in the MATLAB pipeline previous to this thesis [102]. This heavily zero-padded length was found to perform best in the peak detection stages.

¹¹The input impedance of the ADC's channels.

timing. Fortunately, CUFFT FFTs have proven to be much more time-consuming than custom-made functions. The comparison of different FFT libraries [149] could provide useful results and a faster alternative. The reduction of FFT time may also be achieved by adding an analog downconversion stage before the ADC, which would allow removing the first FFT, the IFFT, the digital filtering, and the frequency shift (in total, 68% of current GPU time).

As for custom-made kernels, the slowest is kernel 4, the one responsible for spectrum computation. This occurs even though the signal is already downsampled by the time it reaches this kernel, and it is due to the type of operations performed. It executes logarithms, modules, and complex sums, which are slower than elemental arithmetic operations, such as sums and multiplications. Developing an alternative to the peak detection algorithm that could operate on the time domain [150] would remove this kernel and speedup processing.

Kernels 3 (frequency shift), 6 (filtering), and 7 (calibration) all register similar times. This is once again explained by the type of operations performed. As has already been mentioned, the three kernels execute similar multiplications. The calibration kernel registers slightly lower times because the signal is not zero-padded when it reaches that stage, contrary to the filtering kernel. Even though the signal is also not zero-padded in the frequency shift, this kernel performs a more complicated indexing (involving division operations) than the filtering or calibration kernels.

Comparing between the three different peak detection algorithms, we can see that kernel 10 (maximum peaks detection) has higher time than kernel 8 (external layer peaks detection) or kernel 9 (internal layer peaks detection). Kernels 8 and 9 have very similar workings (the only difference being the direction in which they loop through the signal's samples), so it is to be expected that they register equal timing. Kernel 10 has a more complex loop, because it needs to temporally store and compare the dBm measurement of the samples it loops through, so it is expected that it registers a slightly higher time.

It is also to be expected that kernel 11 (downsampling) registers such a low time. The main operations in this kernel are only load/store calls (as it re-indexes the signal), which are also done in all other kernels, along with their respective main operations. Kernels 12 (polynomial) and 13 (sinusoids) were also expected to register low times. In this case it is due to the fact that only one call of these kernels is needed per image.

Finally, another detailed comparison between the kernels may be seen in Fig. 2.17. In this figure the GPU resource utilization in each kernel is presented. The two types of resources considered are GPU Function Units (which are charged with performing arithmetic operations) and device memory. Function Units can be of various classes (*Control-Flow*, *Double*, *Single*, *Load/Store*, *Special*, *Texture*), depending on the operations they perform. In the graphs presented, the utilization level of Function Units corresponds solely to the class most used (units of other classes register lower utilization levels).

Notice that in this case FFT performance is the same for all instances, regardless of the transform's direction or their position in the processing pipeline. The same occurs for the different peak detection algorithms. As for Function Units, in both cases this occurs because the type of operations performed are always the same, despite the variations of the kernel (such as FFT direction, or layer to be detected). As for Device Memory, the peak detection algorithm needs the same amount of GPU memory in each variation (enough to hold the samples of one ramp per thread), and the slight FFT length variations do not seem to significantly impact its memory use.

As can be seen, the spectrum kernel registers the highest utilization of Function Units, and once again this is due to the intensive operations performed. It is on the verge of being resource-limited. On the other hand, it has low memory usage, equal to kernels 6 and 8. This is due to the fact that these kernels operate on the downsampled signal, so less samples have to be stored in GPU memory. Once again, the filtering and calibration kernel

Figure 2.17: GPU resource use for all kernels, averaged over 100 images. Only one peak detection algorithm is chosen per image. FFTs are done in batched mode with CUDA’s CUFFT library, while all other kernels are custom-made. Profiling is done with NVIDIA’s Visual Profiler.

register equal performance due to their similar workings. Memory usage is the highest in these kernels, along with the frequency shift, because two complete signals are being stored in device memory (the RX signal and the reference used).

An additional design consideration regards thread occupancy on the GPU. The amount of threads necessary in each kernel depends on the number of LFM ramps sent in the ADC data block, and it is usually not a multiple of the warp size (32 threads). To optimize thread scheduling, grid and block sizes have been set to the next multiple of the warp size. This means that some excess threads register as inactive (because they have no assigned operations) on Visual Profiler, which is not only wasteful of GPU resources, but also causes thread divergence in the kernel. A proper study of grid dimensioning is warranted.

2.2.1.2 CPU-code Profiling

Stepping back into Fig. 2.13, it can be noticed that surface reconstruction and image display are not included in the GPU section. These tasks are performed in CPU-executed C++ code using the Visualization Toolkit (VTK) [151]. This is due to the inability of fully adapting the few available open-source CUDA 3-D Delaunay reconstructions to our need. To mitigate this restriction, the Pthreads library is used to parallelize both tasks in CPU hardware (in a multi-core fashion) [152]. A proof of concept image comparison between a previous serial algorithm [102] (with MATLAB-based display) and the new parallel scheme (with VTK-based display) may be seen in Fig. 2.18.

A pool of CPU threads is created, so that as the software is, for instance, using one thread to display an image, it can use a different thread to reconstruct the next. One of the reasons this mechanism was selected is because VTK display windows block the thread that opened them until they are closed, and for this application image refreshing must be done on a single window. Also, due to the use of a pool of threads, even if the time required for the VTK code is higher than the GPU-based code, the user will perceive image refresh at the speed dictated by the signal processing code (with a slight delay only in the first image displayed).

Image refreshing is performed using a VTK timer, programmed for 1 ms interruptions. The timer wakes up the blocked window thread, which checks a flag that indicates whether

or not new data must be displayed. This flag is activated by the other threads in the pool as they finish processing (x,y,z) coordinates (the result of the GPU-executed code) into images. This arrangement and the interaction between all threads in the system is shown in Fig. 2.19.

Figure 2.18: Point cloud image of a mannequin with a simulated explosive device in the abdomen area, in VTK (left) and MATLAB (right).

Figure 2.19: Multi-thread image refresh structure using VTK. Separate CPU threads are assigned to data transfers and image display in order to enable real-time scanning and display.

To profile the CPU section of the code, new configurations are run, each encompassing 100 VTK thread calls (and thus, 100 possible images to be displayed). These configurations are not charged with timing the GPU-based code; the profile starts as soon as (x,y,z) coordinates are available. The results can be seen in Fig. 2.20(a)-(d). These configurations do not consider variations of signal processing nature (such as the type of range layers desired or whether or not to calibrate), because these do not deeply influence the display time of the image. The differences considered regard whether or not a reconstruction of the point cloud is desired, how many threads are set up in the pool, and how many simultaneous range layers to display.

Thread availability goes from 1 free thread to 4, since the CPU provides only 4 cores¹².

(a) Reconstructed surface with 1 layer.

(b) Point cloud with 1 layer.

(c) Reconstructed surface with 2 layers.

(d) Point cloud with 2 layers.

Figure 2.20: Histogram of recorded display times for images, measured over 100 images, with varying number of CPU threads available.

The configurations in Fig. 2.20(a) are of reconstructed surface images, showing only one range layer (internal), and the variation lies on the number of CPU threads in the pool. The number of threads is given in the legend. It can be noticed that all configurations register similar time distributions; thread availability in the pool is not a factor because most executions register times lower than 0.25 s. This means that VTK code needs less time than GPU code, so there is no queue of threads waiting to display an image.

The configurations in Fig. 2.20(b) follow the same structure as Fig. 2.20(a), but for point cloud images. Once again, thread availability is not an issue, because the registered VTK times are lower than the GPU ones. In this case, times are much lower than those in Fig. 2.20(a) due to the removal of the Delaunay-based reconstruction. In fact, most threads register times so low that the C++ functions used to profile the code have rounded the results to 0 s. The required time disparity between Fig. 2.20(a) and (b) poses the question of whether to compromise subjective image quality (a reconstructed surface can highlight threats better than a point cloud) for image refresh (point clouds are more than 10 times faster than reconstructions).

In Fig. 2.20(c), the configurations are once again of reconstructed surfaces, but this time two range layers (external and internal) are displayed. Registered times are higher than in Fig. 2.20(a) because displaying two layers requires additional processing. Once more,

¹²Hyper-Threading techniques have not been considered.

configurations register similar distributions, with the exception of the first configuration (1 thread). This configuration registers half of its images (~ 50) with 0 s. In this case, this is not a rounding effect. It is due to the fact that in this configuration only one CPU thread is available for the VTK code, and display requires more than 0.3 s. This is ~ 0.05 s more than the previously executed GPU code, so one of every two images will find no unoccupied threads in the pool, and will have to be discarded. With just one more thread available (second configuration; 2 threads) this obstacle is overcome (as all images find an available thread to display them), but enabling more threads than necessary (3-4) offers no improvement.

Lastly, in Fig. 2.20(d), the configurations are of point cloud images, showing two range layers (external and internal). Once again, point cloud times are much lower than those of reconstructed surfaces, but comparing with Fig. 2.20(b), we can see that times have risen slightly. As mentioned earlier, this is because displaying two layers requires slightly more code execution.

The bimodal behavior of the distributions, noticeable specially in Fig. 2.20(a) and (b), is due to the fact that although 100 images were represented, this was done by alternating between the same 2 point clouds. Although both radar scans are made up of 6000 possible spots, the actual number of spots that surpass the power threshold (2366 for one, and 3245 for the other), and the surface reconstruction performed among them, vary from image to image. These effects produce slight differences in timing; for instance, further reducing the number of points by half in both images lowers their average by 0.1 s (but also deteriorates reconstruction quality).

In any case, thanks to the thread pool mechanism 4 fps (limited by GPU processing and transfers) can be guaranteed for all images, although it must be noted that except configurations in Fig. 2.20(c), all others register VTK times that would allow faster refresh. Nevertheless, the use of the VTK code has introduced additional GPU-to-CPU memory transfers and control functions that have unnecessarily bloated time per image. The programming of Delaunay reconstructions and image display on CUDA code is a top priority to further speed-up processing [153].

2.2.1.3 Complete Profiling

Finally, the total processing and display time (considering both GPU and VTK code) was profiled on 12 new configurations, each displaying 100 images. The averaged results may be seen in Table 2.6. As can be noticed, the total time required is not exactly the same as the GPU processing time (Table 2.4), but slightly higher, possibly due to previously unprofiled control functions. According to this profile, in the worst average scenario (0.27301 s) the refresh would be 3.66 fps.

2.2.2 Improved Processing Architecture

Summarizing the results from Section 2.2.1, two main limitations for faster image refresh can be recognized:

- **FFTs.** The FFTs in the algorithm account for more than 60% of the total GPU processing time. A reduction of the processing time per FFT, or the total number of transforms, would produce a substantial improvement in refresh speed.
- **Hardware.** The GPU profiled in the previous section is a GeForce GTX 970, which is a mid-tier consumer-level unit released in 2014. GPUs with additional processors, arithmetic units, and memory size have since then been commercially available. More

Table 2.6: Time per image in 100-image configurations.

Config.	N ^o of peaks	Range layer chosen	Reconstruction	Time (s)
1	1	External	No (point cloud)	0.25418
2	1	Internal	No (point cloud)	0.25214
3	1	Max. peaks	No (point cloud)	0.25486
4	1	External	Yes	0.25525
5	1	Internal	Yes	0.25839
6	1	Max. peaks	Yes	0.25784
7	2	Ext. + Max.	No (point cloud)	0.26366
8	2	Ext. + Int.	No (cloud)	0.25981
9	2	Int. + Max.	No (point cloud)	0.26108
10	2	Ext. + Max.	Yes	0.27301
11	2	Ext. + Int.	Yes	0.26408
12	2	Int. + Max.	Yes	0.26677

importantly, the CPU used is a low-grade system (ASUS P8Z68-V LX motherboard with an Intel Core i5-2500K processor, both released in 2011), whose memory bandwidth limits the CPU-to-GPU and GPU-to-CPU data transfers.

These restrictions have been overcome by re-designing the signal processing algorithm, and implementing it in a different hardware platform. To reduce the number of FFTs executed per frequency ramp, the downsampling stage has been eliminated¹³. This allows removal of the FFT performed before downsampling and the IFFT, but also the low-pass filtering, frequency shift, and sinusoid kernels. A diagram of the updated signal processing chain may be seen in Fig. 2.21. Along with this removal, the processor has been updated to an Intel Core i7-4790 (released in 2014) and the GPU to a GeForce TITAN X (released in 2016).

Figure 2.21: In black, main stages of the improved signal processing algorithm, in execution order from left to right. In red, stages removed from the previous architecture in Fig. 2.6.

Additionally, the spectrum computation kernel (the slowest custom-made function in the algorithm) has been optimized by eliminating the logarithm calculation¹⁴ and the constant summand in (2.2.1). The resulting spectrum is then calculated simply as

¹³The length of the final, and now only, FFT changes from 4096 to 8192. This value is the next power of two of the chirp's sample length. The 65536 standard FFT length is not used to reduce FFT execution time, at the possible cost of peak detection precision.

¹⁴Thus operating now with linear units.

$$S_{d_k} = |\mathcal{F}(s_d[n])|, \forall n, k \mid 1 \leq n \leq N \wedge 1 \leq k \leq \frac{N}{2}, S_{d_k} \in \mathbb{R} \quad (2.12)$$

2.2.2.1 GPU-code Profiling

As in Fig. 2.15, the GPU processing time required for 100 images of 6000 possible chirp periods each has been profiled. The results may be seen in Fig. 2.22, and the updated average times are presented in Table 2.7.

Figure 2.22: Average GPU-based signal processing rate per 6000-LFM-ramp image for various ADC data sizes. Configurations vary depending on the signal processing and display choices seen in Table 2.7. Each configuration time is averaged over 100 images, per datapoint in the x -axis. (a) Shows the required GPU processing time per image. (b) Shows the processing speed in images per second [the inverse of (a)].

Table 2.7: CUDA signal processing configurations

Config.	Peaks	Layer	Calibration	Time (s)*	Reduction [†]
1	1	External	Yes	0.1051	57.41%
2	1	External	No (windowing only)	0.1058	57.29%
3	1	Max. peaks	Yes	0.1041	57.96%
4	1	Max. peaks	No (windowing only)	0.1047	57.87%
5	1	Internal	Yes	0.1057	57.33%
6	1	Internal	No (windowing only)	0.1062	57.11%
7	2	Ext. and Max.	Yes	0.1042	58.29%
8	2	Ext. and Max.	No (windowing only)	0.1045	58.10%
9	2	Ext. and Int.	Yes	0.1060	57.41%
10	2	Ext. and Int.	No (windowing only)	0.1056	57.69%
11	2	Int. and Max.	Yes	0.1054	57.89%
12	2	Int. and Max.	No (windowing only)	0.1045	58.25%

*Per image, averaged over 100 images, and using 3000-chirp blocks.

[†]Reductions in time compared to Table 2.4.

The asymptotic trend of the processing time is still present after modifications, indicating that the occupancy problem is not GPU-model-dependent. Regardless, processing times have improved in average by $\sim 57\%$, resulting in an image refresh rate greater than 9 fps. Time

reductions are in the order expected from Section 2.2.1, where all FFTs accounted for $\sim 63\%$ of processing time. Note that in the improved algorithm a batch of FFTs, before peak detection, still remains and its size is larger than before due to the removal of the downsampling.

Individual kernels have been profiled, with the results seen in Fig. 2.23 and Table 2.8. As with the original algorithm, the profile is performed with NVIDIA Visual Profiler [148], and the results are once again averaged over 100 images. The sum of kernels 1-3, 7, and one of 4-6 would be the total processing time required per average image. As can be seen, all kernels have benefited from reductions in time greater than 30% due to the faster clock rate (20.3% quicker) and memory bandwidth (87.5% quicker) of the new GPU. In the worst-case scenario is selected (choosing the slowest of the three peak detection algorithms), the total kernel time per image is just 0.02165 s, which is ~ 0.083 s less than the results seen in Table 2.7. Once again, this difference is partly due to memory transfers, which now account for 0.046 s (55.42% of the time-per-image difference between Fig. 2.23 and Table 2.7). A redesign of the signal processing algorithm, focused on more efficient transfers and the use of multiple CUDA streams, could further improve refresh rate. Nonetheless, considering only GPU-executed kernels (0.02165 s per image), refresh could now be as fast as 46.19 fps.

Figure 2.23: GPU processing time per image for the remaining kernels, averaged over 100 images. Only one peak detection algorithm is chosen per image. FFTs are done in batched mode with CUDA’s CUFFT library, while all other kernels are custom-made. Profiling is done with NVIDIA’s Visual Profiler.

Table 2.8: Distribution of image time over kernels

Number	Kernel	Time (ms)	Time (%) [†]	Improvement (%) [‡]
1	FFT	16.09	74.3	31.85
2	Calibration	3.03	14.0	64.27
3	Spectrum computation	2.15	10.0	84.26
4	Peak (maximums)*	0.37	1.7	94.14
5	Peak (external)*	0.23	1.1	94.64
6	Peak (internal)*	0.23	1.1	93.21
7	Polynomial	0.01	0.0	75

[†]Percentage over sum of total GPU kernel time for an image.

[‡]Reductions in time compared to Table 2.5.

*Only one of these kernels is chosen per image.

In the new algorithm, the remaining FFT accounts for 74.3% of the total GPU processing (transfers not included) time per image. Of that 74.3%, the distribution between FFT plan-

ning and execution would now be 26.5% and 47.8%, respectively. Relatively, the decrease in processing time for the FFT has not improved as much as other kernels ($\sim 30\%$ improvement vs. more than 64% improvement for other kernels). This is due to its increased length after removing the downsampling stage.

The slowest custom-made GPU function is now kernel 2, responsible for calibration, due to the optimization of the spectrum computation kernel. It is now slower than kernel 3 because of the removal of the logarithmic function. Computing the spectrum in logarithmic units in kernel 3 is enough to double its registered time and surpass the calibration kernel.

The peak detection kernels (4-6) were expected to record low times. No arithmetic operations are performed in these stages, and the data samples to process are reduced, due to the symmetry of the Fourier spectrum of real signals and the range gating. Comparing between the three algorithms, kernel 4 (maximum peaks) has higher time than kernel 5 (external layer peaks), which is equal to kernel 6 (internal layer peaks), as was the case in Section 2.2.1. The large decrease in processing time ($\sim 94\%$) for the peak detection kernels is due to the enhanced effect of improving memory bandwidth (because of the multiple read-and-store operations while sorting samples).

The resource utilization in each kernel is presented in Fig. 2.24. The two types of resources considered are once again GPU Function Units and memory bandwidth. Notice that, as in the previous algorithm, the peak detection kernels maintain their performance regardless of the type of layer selected.

Figure 2.24: GPU resource use for all kernels, averaged over 100 images. Only one peak detection algorithm is chosen per image. FFTs are done in batched mode with CUDA’s CUFFT library, while all other kernels are custom-made. Profiling is done with NVIDIA’s Visual Profiler.

As can be seen, the FFT and the spectrum computation are the two kernels that use the most Function Units. This is to be expected, since the remaining kernels perform very basic arithmetic operations (of which the GPU has more units available). The FFT has increased from $\sim 80\%$ Function Unit utilization in the previous algorithm to 95% utilization now, in spite of being implemented in a higher-performance GPU. This is most likely due to the increase in FFT length and the removal of other kernels, which requires and allows it, respectively, to use additional units. The performance of the other kernels has been more or less maintained. It is interesting to note that if the *log* function were added to kernel 3, another 5% in Function Unit utilization would be increased to this stage, surpassing the FFT and fully limiting the kernel.

As for device memory, the high usage of the calibration kernel can be explained by the

storing of the reference file that calibrates the IF signal, while the usage of the other custom-made kernels is limited to storing only IF signal samples from half of its spectrum. The FFT has additional storage instructions during its planning stage to prepare for its execution, which produces a slight increase of memory usage with regard to kernels 3 - 7.

2.2.2.2 CPU-code Profiling

The image display section of the code, executed on CPU, has not been modified from the previous design. Nonetheless, since the full algorithm is run on a pipeline scheme where the GPU code executes first, the image refresh results have varied. Once again, to profile this section of code, 100 images were profiled after (x,y,z) coordinates are available. The results can be seen in Fig. 2.25(a)-(d). Although the CPU has changed, thread availability still goes from 1 to 4 free threads.

(a) Reconstructed surface with 1 layer.

(b) Point cloud with 1 layer.

(c) Reconstructed surface with 2 layers.

(d) Point cloud with 2 layers.

Figure 2.25: Histogram of recorded display times for images, measured over 100 images, with varying number of CPU threads available.

Fig. 2.25(a) shows processing times similar to Fig. 2.20(a). At most, an average of ~ 0.01 s has been deducted from the processing time of reconstructed 1-layer images. This would indicate that the reconstruction process via Delaunay triangulation has not sped up as a result of the CPU change, highlighting the importance of porting this algorithm to GPU. Another interesting result in this graph is the number of images ($\sim 50\%$) registering 0 s of processing time. As was the case in Fig. 2.20(c), this is due to the fact that GPU processing now requires just ~ 0.105 s per image, while VTK code still requires between 0.15 s and 0.25 s to display it. The display time is approximately double the processing time, meaning that

1 in every 2 images has no free thread available when only 1 thread is allocated in the pool. Thus, half of the images cannot be displayed. If 2 or more threads are reserved, this situation is solved.

As already happened for the initial design in Fig. 2.20(b) and Fig. 2.20(d), the point cloud representations of Fig. 2.25(b) and Fig. 2.25(d) show no problem in matching GPU processing speed. Practically all images register display times rounded to 0 s due to the fact that these take almost no time to represent.

Lastly, in Fig. 2.25(c), the problem with Delaunay reconstruction over CPU is still present. Since in this case the images have 2 layers, reconstruction takes double the time from Fig. 2.25(a), meaning that now display threads can take as much as ~ 0.4 s. Since GPU processing is ~ 0.105 s per image, only when 4 threads are available in the pool no images are lost. For any other number of available threads (1-3) at least 25% of the images register 0 s of display time, meaning they have not been represented. As for future modifications, this situation puts emphasis on the need for migrating the image display code to GPU, because, since no more threads can be allocated in the CPU used, the VTK-executed software will not be able to keep up with further improvements on the GPU processing section.

2.2.2.3 Complete Profiling

Nevertheless, due to the thread pool scheme, image refresh is limited simply by GPU processing. According to Table 2.7 the worst scenario would register 0.1062 s of processing time, equaling a 9.42 fps image refresh. The complete processing and display time has also been profiled to provide the final benchmark of image refresh. The results may be seen in Table 2.9, where the worst average is 0.1070 s, resulting in 9.35 fps of image refresh.

Table 2.9: Time per image in 100-image configurations.

Config.	N ^o of peaks	Range layer chosen	Reconstruction	Time (s)
1	1	External	No (cloud)	0.1056
2	1	Internal	No (cloud)	0.1061
3	1	Max. peaks	No (cloud)	0.1044
4	1	External	Yes	0.1054
5	1	Internal	Yes	0.1061
6	1	Max. peaks	Yes	0.1070
7	2	Ext. + Max.	No (cloud)	0.1048
8	2	Ext. + Int.	No (cloud)	0.1049
9	2	Int. + Max.	No (cloud)	0.1063
10	2	Ext. + Max.	Yes	0.1064
11	2	Ext. + Int.	Yes	0.1060
12	2	Int. + Max.	Yes	0.1052

Nonetheless, it is also fair to recall from the GPU profiling that ~ 46 fps refresh is readily attainable with the CUDA algorithm (Fig. 2.23), if memory transfers are performed parallel to execution and two more CPU threads were available. This is illustrated in Fig. 2.26, where asynchronous with two CUDA streams is used. This allows bypassing delays due to memory transfers, because while an image is being processed in one CUDA stream, the data from the next image is being transferred with the other CUDA stream. Of course, to be able to do this, all ADC samples of an image need to be available at least as fast as the memory transfers, meaning that the scanning mechanism would need to be speeded up. Because GPU-processed data would then be available faster (at a rate of 0.022 s per image, due to two processing blocks of 0.011 s each), and assuming that VTK surface reconstructions

require in average ~ 0.3 s, the CPU would need at least six threads to process and display the images without losing any frames. Without taking into account the initial delay until the first image is displayed, the refresh rate would then be set by the time difference between the end of VTK processing for consecutive images, which due to the pipeline scheme is equal to the total GPU processing time (0.022 s per image, or ~ 46 fps).

Figure 2.26: Possible timeline of image processing and display, with parallelized GPU processing and host-to-device memory transfers with the use of two different CUDA streams. ADC data samples of all images must be at least synchronized with GPU transfers and at least 6 CPU threads are necessary.

For summarization purposes, the final improvement compared with the MATLAB implementation may be seen in Table 2.10.

Table 2.10: Kernel timing comparison between the final GPU algorithm and a MATLAB-based implementation

Kernel	CPU-GPU (ms)	MATLAB (s)	Improvement Factor
FFT	16.09	118.89113	7389
Calibration	3.03	0.20224	67
Spectrum computation	2.15	0.30573	142
Peak detection (max.)	0.37	12.51003	33811
Peak detection (ext.)	0.23	12.42096	54004
Peak detection (int.)	0.23	12.62477	54890
Polynomial	0.01	0.03767	3767
Point cloud display	~ 1	0.64487	645
Reconstruction	~ 150	1.34448	8963

Some conclusions can be inferred from the results of the improved design. These are summarized as follows:

- The use of the TITAN X GPU and the more modern CPU are highly responsible for the increment in image refresh. The improved algorithm by itself (without the updated CPU and GPU) would produce only 6.56 fps. If just the CPU is changed, that figure rises to 8.38 fps, primarily due to the acceleration of CPU-to-GPU memory transfers

from 84 ms per image to 46 ms per image (due to an increase in memory of bandwidth from 6.874 GB/s to 12.383 GB/s). If the GPU is also changed, the frame rate rises to the final 9.35 fps refresh. This seems to indicate that further modifications to the hardware, such as newer GPUs/CPUs, could produce quicker rates without algorithm modifications.

- The GPU's full potential has not yet been tapped with this algorithm. Some of the algorithm's main limitations are from the Delaunay reconstruction and other image processing and display functions, which are graphical applications, for which GPUs are specifically designed. An implementation of these functions on GPU would remove these limitations. Additionally, some areas of the processing algorithm that are already run on GPU could benefit from further parallelization. For instance, each LFM ramp is assigned a single GPU thread during peak detection, but parallel reduce techniques¹⁵ could provide threads for individual signal samples.
- The utilization of Function Units is very high even in kernels that perform simple operations (such as the calibration kernel, which presents $\sim 80\%$ usage). This is because the kernels are programmed to deal with few data types, which means that most operations can only be performed by a single class of Function Units. This high usage can be a bottleneck if asynchronous processing is implemented (such as in Fig. 2.26), since few resources are available to other concurrent kernels. Diversifying data types and reducing the number of floating point operations could produce additional improvements.

2.2.3 System Performance

The signal and image processing architecture presented in the previous subsection has been implemented in the 300 GHz imaging radar. This architecture has allowed capturing real-time video sequences of moving targets during on-field measurements. Although the mechanical scanning still limits the image refresh to ~ 2 fps, the software now permits continuously viewing the scene at that speed.

The final real-time system also takes into account low-frequency electronics and mechanical aspects to synchronize the signal processing stages with the scanning movement. In our implementation this is done via a custom-made printed circuit board (PCB) that manages the outputs of the scanning motors' encoders, providing real-time data regarding their positioning. These data are sent to the PC via Ethernet connection and are then provided to the image processing architecture. In addition, data transfer limits in the ADC used restrict the sampled throughput. This has forced the sampling strategy to be switched from full-band (corresponding to ranges from 0 m to ~ 16.7 m) to bandpass (corresponding to a range window of 30 cm around the standoff distance). Bandpass sampling allows decreasing the sampling frequency and thus the throughput of sampled data. To do this without creating an aliasing effect, an analog bandpass filter is added to the radar's IF path before sampling. All of these considerations are described in detail in Annex C,

Radar imaging has been tested with people-scanning scenarios in a controlled laboratory environment, with a standoff set-up shown in Fig. 2.27(a). Since the threat-detection capabilities of the radar have already been demonstrated in previous publications [102, 104, 123], the scenarios presented here are focused on validating the real-time scanning performance through the capture of subject movement. A static object (metallic frame holding a thin bar with a sphere attached), seen in Fig. 2.27(b), is placed in the background to confirm range and power stability for non-moving objects, as well as provide contrast with the moving subject.

¹⁵Of which there are examples in the literature [154].

Figure 2.27: Photograph of (a) the real-time scanning set-up for a static metallic object, and (b) the background target used (a rectangular metal frame holding a thin bar with a sphere attached at one end).

The first scenario deals with a male entering and exiting the FoV, while extending and closing his arms. Sample images from a sequence captured at ~ 1 fps can be seen in Fig. 2.28.

The second scenario depicts the same male entering and exiting the FoV, while crossing his arms in front of him. The static metal frame from Fig. 2.27(b) is still present in the background. This time, the radar has been configured to represent range information. Sample images from a sequence captured at ~ 1 fps can be seen in Fig. 2.29.

The last scenario depicts the male entering the FoV, with his arms extended. The radar is still configured to represent range information, and the metallic frame from Fig. 2.27(b) is present. This time the sequence has been captured at ~ 2 fps (limited by motor speed), which allows better movement resolution, but produces less dense point clouds (specially noticed in the metallic frame). Sample images are shown in Fig. 2.30.

Figure 2.28: Sample images from a continuous video sequence at ~ 1 fps of a male entering and exiting the FoV (while extending and closing his arms), as captured by the radar. The static metal frame from Fig. 2.27(b) is present in the background. The chronological sequence starts from the top left corner, left to right, and is also indicated by the numerical order attached to each image. Colorbar: Received power in dBm.

Figure 2.29: Sample images from a continuous video sequence at ~ 1 fps of a male entering and exiting the FoV (while crossing his arms in front), as captured by the radar. The static metal frame from Fig. 2.27(b) is present in the background. The chronological sequence starts from the top left corner, left to right, and is also indicated by the numerical order attached to each image. Colorbar: Range in meters.

Figure 2.30: Sample images from a continuous video sequence at ~ 2 fps of a male entering the FoV (with his arms extended), as captured by the radar. The static metal frame from Fig. 2.27(b) is present in the background. The chronological sequence starts from the top left corner, left to right, and is also indicated by the numerical order attached to each image. Colorbar: Range in meters.

2.3 Millimeter-wave Scattering and Multistatic Configuration

It has already been stated how, in the past decades, monostatic security screening radars have progressed into feasible threat detection prototypes. Although these systems have proven to be successful in revealing concealed weapons in static and controlled environments, future security checkpoints demand high detection probabilities even with moving and noncooperative subjects. In this regard, monostatic systems can only generate images from a single viewpoint, omitting the detection of person-borne threats in non-viewed areas of the body. For these reasons, active imagers, which have traditionally been implemented in single-transmitter monostatic architectures [102, 126, 136], are recently progressing towards multi-transmitter [92, 135] and multistatic [155] configurations. The main purpose of this evolution is to adapt systems to more realistic imaging scenarios where multiple viewpoints are needed to accurately represent and detect a threat, as requested by stakeholders such as the DHS [98] or the TSA [99].

These changes in architecture also serve as an initiation into future on-the-move walk-through configurations, which allow for faster passenger throughput and are easily adaptable to airport infrastructures [139, 156]. Additionally, in some cases, multistatic architectures have proven to achieve increased cross-range resolution in their images, when compared to traditional monostatic configurations [157]. In this regard, the diagrams in Fig. 2.31 illustrate some of the common imaging configurations for security screening.

Figure 2.31: Illustrations of different imaging architectures according to the number of radiating elements and their placement. (a) Shows a standoff monostatic case, (b) shows walkthrough panel arrays, and (c) shows a standoff multistatic scenario.

In Fig. 2.31(a) a conventional standoff monostatic architecture is shown, where the FoV is limited to a small area and to the viewpoint of the single sensor. An approach to increase the FoV and allow multiview imaging could be the modification of scene scanning. A technique for adaptable range focusing, which can continuously provide images as the subject moves towards the radar, has already been presented [158]. Nevertheless, this method can only compensate movements in the range direction. A similar approach would be the circular

movement of the radar around the subject, as done with commercial success in L3's ProVision systems [101], although this mechanism would require a static target in a specific pose throughout the scanning. A more promising technique is the tilting of the main reflector, such as in [159], to move the FoV as the subject walks past the radar, although it would only allow imaging the person from a single side.

A solution to image multiple sides of the target without disturbing the subject's flow would be the use of parallel walkthrough panel arrays of TRXs [139, 156], such as shown in Fig. 2.31(b). Some of the disadvantages of this architecture in comparison to standoff imaging would be the higher cost and complexity due to the large number of TRXs required, as well as the close proximity of the target.

A simple approach to achieve both multiview imaging and unperturbed subject flow would be placing a set of standoff radars around the walkway, as shown in Fig. 2.31(c). In this scheme, independent FoVs for each reflector system are obtained, as well as bistatic propagation paths present between each pair of radars, further increasing image information. Of course, the main detriment of this architecture in the mm-wave and THz bands would be the large cost associated to duplicating RF equipment, specially costly reflector systems.

Another important advantage in multistatic imaging configurations in high-frequency bands is the possibility to avoid some of the adverse effects of specular scattering at these frequencies. In specular scattering most of a target's reflections are directed at a single angle, whereas in diffuse scattering these reflections are spread out over a wider range of directions. The angle at which the waves are reflected depends mainly on the angle of incidence and the target's surface. In rough surfaces where micro-curvatures have comparable size to the transmitted wavelength the scattering is primarily diffuse, whereas in smooth facets (such as those in many metallic threats like knives or guns) the reflections are primarily specular in nature. Fig. 2.32 includes diagrams of these two types of reflections.

Figure 2.32: (a) Specular reflection from a smooth surface, where $\theta_i = \theta_r$ (b) Diffuse reflection from a rough surface, where $\Delta\theta_r$ contains various reflection angles.

In an imaging radar using a focal system to scan the target, specular scattering can produce critical imaging artifacts [22]. When the return angle is aligned with the radar's orientation, the scattered signal will be fully captured by the radar, and therefore a powerful glints and bright speckle will be present in the resulting image, if returned power data are represented. If the scattering angle does not coincide with the radar's orientation, little to no returned power will be received by the radar, and areas of the FoV will vanish below the radar's detection threshold. These effects are even more significant in standoff imagers, where the larger distance to the target implies worse angular tolerance under random alignments

[160]. Because glints and speckles still contain useful information regarding the target, and image contrast can be compensated [161], vanishing can be considered the most critical effect.

Although multiple simulated and measured examples of target vanishing in mm-wave imaging radars have been presented [162–164], the effects of specular reflection as a person moves through the checkpoint are still an unsolved problem. Nonetheless, because specular scattering depends on the incidence angle of the transmitted radiation, this vanishing effect can be averted through the reorientation or repositioning of the radar. Furthermore, if instead of repositioning the radar, multiple viewpoints of the target are simultaneously available, as in a multistatic architecture, the vanishing in each individual image can be compensated with results from the others.

In this text, a multistatic frequency-modulated continuous-wave (FMCW) imaging radar configuration is evaluated, based on a multi-bistatic configuration shown in Fig. 2.33. This configuration is a cost-optimized and simplified alternative to the multistatic scheme in Fig. 2.31(c) that offers multiview imaging while also managing cost concerns from replicating RF equipment. In the proposed architecture, a single radar is used to receive focused radiation, while multiple independent TXs are placed around the walkway. Each TX can then provide the radar with a bistatic path to enhance its monostatic image. With optimal placement and mechanical steering of the separated TXs, a large section, if not all, of the subject's body could be imaged as it moves.

Figure 2.33: Standoff multi-bistatic imaging configuration, simplified from the multistatic architecture in Fig. 2.31(d).

Fig. 2.34 presents a block diagram of our proof-of-concept implementation. In this architecture, the 300 GHz radar's RX chain is used as TRX, and placed in the reflector system. At the same time, the radar's TX chain is placed at an independent location (away from the reflector system) globally illuminating the same FoV, and providing a wider observation angle than the monostatic case.

Figure 2.34: Block diagram of the implemented multi-bistatic architecture using an independent TX. To generate multiple bistatic images, the TX can be moved throughout the scene.

Due to the TRX's ability to function simultaneously as a TX and RX [123], the TX chain is not necessary to generate monostatic images, and thus, it can be separated from the reflector system and used to illuminate the scene. This enables using the TX as an independent

illuminator, and therefore, one bistatic and one monostatic image are simultaneously possible. The RX side of the TRX is able to simultaneously acquire the reflections from its own radiation and the ones from the independent TX, producing the monostatic and bistatic images at the same time.

Through placement of the TX at different locations, a multistatic image of a static target can then be generated by combining all the resulting monostatic and bistatic images during offline processing. The angular diversity achieved allows to maintain target detection in at least one of the resulting images. To place the TX in its separated location, a conventional photography tripod with a 3-D printed mount is used, as seen in Fig. 2.35. The tripod's head presents a 5-axis movement, allowing selection of specific pitch, yaw and roll angles for the TX. To illuminate the scene, the TX uses its COTS horn antenna, with a 12° beamwidth¹⁶.

Figure 2.35: Separation of the 300 GHz TX chain from the reflector system with a photography tripod for multistatic measurements.

To ensure system coherence, the TRX and TX share the radar's microwave chirp source, which is later multiplied in each chain to the transmitted mm-wave frequency range. The signal is distributed at microwave frequencies to avoid the higher losses and lack of flexible components in the mm-wave band, at the cost of needing more than one RF upconversion chain. To distribute the chirp to the TX chain, a low-loss coaxial cable¹⁷ is used at the output of the radar's microwave subsystem. The coaxial used is 3 m long in order to provide enough separation from the reflector system to achieve meaningful angular diversity. Although the cable is a low-loss model, because of its length and the chirp's high frequency, ~ 6 dB losses have still been observed in the 16 GHz distributed signal.

In the following subsections, simulated and measured imaging results, from monostatic, bistatic, and multistatic viewpoints, are presented to show the usefulness of such a configuration in scenarios dealing with common threats.

2.3.1 Millimeter-wave Scattering Behavior

In the mm-wave and submm-wave bands wavelengths have comparable size to a target's surface roughness, therefore the scattering behavior of reflected signals presents high variability from one imaging scenario to another. This subsection is dedicated to introducing some imaging effects derived from this behavior, as well as some of the theoretical background for the multi-bistatic FMCW imaging scheme implemented.

¹⁶The TRX also includes a same-model horn, as well as the reflector system.

¹⁷Model PE300-120 from Pasternack.

Figure 2.36: (a) Bistatic FMCW radar configuration with one TRX and one TX. (b) Sawtooth FMCW signals in the bistatic configuration.

Fig. 2.36 illustrates a bistatic range-measurement scenario based on FMCW waveforms. In this scenario a static TRX is focused onto the target, while an independent TX (which can rotate around the target with a bistatic angular diversity of θ degrees) simultaneously illuminates it in a global manner. In this example, BW is the transmitted bandwidth, T_m is the chirp period, f_{bi} is the beat frequency, and Δt_i is a time offset between the transmitted and received signals given by

$$\Delta t_i = \frac{R_i}{c_0} \quad (2.13)$$

where R_i is the distance travelled by the transmitted and reflected signal, and c_0 is the speed of light in vacuum.

In the bistatic scenario in Fig. 2.36, there are two beat frequencies produced in the TRX by mixing a sample of its transmitted signal with the reflections of its own transmission (monostatic path) and of the TX's (bistatic path). The monostatic beat frequency will be proportional to the monostatic distance, while the bistatic beat frequency will be proportional to the bistatic distance. The corresponding time offsets that produce these frequencies are

$$\Delta t_i = \begin{cases} \frac{2R_{TRX}}{c_0} & \text{if } i = 1 \\ \frac{R_d}{c_d} + \frac{(R_{TX} + R_{TRX})}{c_0} & \text{if } i = 2 \end{cases} \quad (2.14)$$

where R_{TRX} is the range from the TRX to the target, R_d is the distance corresponding to signal distribution between the chirp source and the separated TX, c_d the propagation velocity in the distribution medium, and R_{TX} is the range from the TX to the target. With beforehand knowledge of the signal's distribution length, R_d is known, and thus, for an unknown target position, R_{TX} can be extracted using the measured R_{TRX} from its corresponding beat frequency. Thus, the monostatic distance will simply be two times the distance between TRX and target, to account for transmission and reflection paths. The bistatic distance will be the sum of the distribution medium's length, the distance between TX and target (transmission path), and the distance between target and TRX (reflection path). In our implementation, the x and y coordinates will be the same for simultaneous monostatic and bistatic images, because they share the same scan pattern from the common reflector system, while their z coordinates will be different due to their different propagation paths. In our implementation, when combining multiple monostatic and bistatic point clouds

into a single multistatic representation all 3-D coordinates from each individual point cloud are taken into account (no filtering or clustering is performed).

Fig. 2.36(b) also assumes a realistic target surface which can cause reflections in a range of beat frequencies Δf_b due to various penetrated range layers. This highlights the importance of an appropriate TX placement to ensure that Δf_{b1} and Δf_{b2} do not overlap in frequency, creating an ambiguity in the TRX when assigning beat frequencies to each transmission. If the beat frequencies reflected from each transmission are separated, two simultaneous range measurements of the same target can then be performed. For an imaging radar case, this means that two simultaneous images of the scene can be generated. To do this, two frequency windows will need to be evaluated during processing to form the separate images, as illustrated in Fig. 2.37 (where the TX frequencies show degradation and non-linearities, as mentioned in Section 2.2). The figure also shows how, in order to properly combine the point clouds into a single image, the z coordinates must be compensated with a range offset (d_{offset}) corresponding to the difference between bistatic and monostatic propagation paths.

Figure 2.37: Diagram of the processing pipeline to obtain and combine simultaneous monostatic and bistatic point clouds. Colorbars: Range in meters.

In our implementation, the selection of the TX's placement is conditioned by multiple variables. For instance, the TX cannot be placed blocking the monostatic path and it also cannot be placed in a distance such that Δt_1 is equal to Δt_2 in order to ensure two separate beat frequencies. Furthermore, to properly illuminate the entire FoV, the TX must be placed far enough from the target to cover the area with its antenna's beamwidth. In addition, the ADC's peak sampling rate will restrict the maximum beat frequencies (and thus ranges) detectable. Lastly, an obvious limit in this specific case is the cable's length, which at 3 m long provides a peak angular diversity from the TRX of $\pm 25^\circ$ around the target.

To illustrate the necessity and benefits of a multistatic imaging configuration, several target scenarios have been analyzed. The first target is a smooth metallic plate scanned with the TRX in monostatic configuration, chosen to highlight the behavior of highly-specular metallic threats. The target is first imaged in an orientation that maximizes the power returned to the TRX, and then also rotated around its vertical axis using a rotational positioner. The measured results can be seen in Fig. 2.38(a).

The maximum returned power from the sheet is recorded as it is being rotated in 0.5° steps around its vertical axis. The power received in a fixed monostatic configuration is compared with the power registered from two TX placements forming $\theta = 6^\circ$ and $\theta = 12^\circ$ bistatic angles with the radar's RX. The results can be seen in Fig. 2.38(c). The highly-specular target presents 61 dB losses at 3° in the monostatic case (after which returned power reaches the radar's monostatic noise floor). Because of the vanishing associated to such losses, this scenario highlights the difference between the radar's effective and maximum FoVs. The maximum FoV is the largest area the radar is capable of imaging, and will be dependent

Figure 2.38: (a) Measurement target consisting of a smooth metallic sheet held with a rotational positioner. (b) EPS foam slab in the same positioner. (c) Returned power from both targets in a monostatic TRX, and from TX placements at $\theta = 6^\circ$ and $\theta = 12^\circ$, as the target is being rotated around its vertical axis. (d) Zoomed-in view of the EPS results, with the metallic data removed for better visibility.

only on the reflector system and scanning mechanism. Meanwhile, the effective FoV can be defined as the area within the maximum FoV where the radar generates detections, and will depend on the detection threshold, the target's scattering properties, and its orientation.

If the rotated plate is also illuminated from an independent TX placed with a $\theta = 6^\circ$ angular diversity, the rotational losses are overcome in the bistatic image which presents maximum return at precisely 3° degrees. In our implementation the bistatic curves present a higher noise floor compared to the monostatic data, due to predominance of phase noise from higher transmit power in the bistatic path, compared to thermal noise in the monostatic path [123]. Of course, the bistatic reflections also suffer from specular effects, leading to equally-large losses if the target continues rotating. The losses in the bistatic image at $\theta = 6^\circ$ when the plate continues rotating can then be bypassed with the placement of another TX at a larger angular diversity, as in the measured case at $\theta = 12^\circ$. To fully solve this problem multiple TXs covering all possible angles in a true multistatic configuration are required.

To now analyze the impact of rotational losses in imaging scenarios with less specular behavior, the measurement set-up in Fig. 2.38(a) has been repeated for a different material. In this new experiment, the target is now a 3-cm-thick block of expanded polystyrene (EPS) foam shown in Fig. 2.38(b), commonly used in packaging. The internal structure of the polymer, which the 300 GHz radiation penetrates [166], is formed from multiple pellets, providing a significant diffuse component to its reflections [167]. The resulting data can be seen in Fig. 2.38(c), and zoomed-in in Fig. 2.38(d). Because of the signal's penetration through the EPS, the bistatic curve displays two power peaks (at 6° and 11°) due to different bistatic angles of the front and back surfaces [167], with the front surface evidently reflecting more power. In both configurations the polymer still exhibits a noticeable specular component to its reflections, as the maximum received power is still located at the rotation angles corresponding to the different radiation sources (0° and 6° for the monostatic and bistatic cases, respectively). Nonetheless, the power losses due to target rotation are much lower than for the highly-specular metallic sheet, indicating that a strong diffuse component is also present. In targets with such a diffuse behavior, the advantages of a multistatic imaging configuration might not be so apparent, as a 6° rotation only implies ~ 5 dB losses, but this attenuation could still prove large enough to fully or partially vanish a weak threat from an image.

Fig. 2.39(a)-(c) show the monostatic imaging results of the metallic sheet for 0° , 3° , and 6° rotations. As can be seen, when the sheet is oriented towards the radar in Fig. 2.39(a), received power is centered on the target. A 3° rotation in Fig. 2.39(b) already reduces the power in the sheet's central area approximately to the radar's detection threshold (-75 dBm)¹⁸, which would cause its vanishing. The maximum peaks [which are decreased ~ 15 dB compared to Fig. 2.39(a)] in the rotated scenario are now found in its left-side edge, because it is the only area oriented towards the TRX. When the sheet is further rotated to 6° in Fig. 2.39(c), the maximum returned power is again decreased another ~ 5 dB. As in Fig. 2.39(a)-(c) for the metallic sheet, Fig. 2.39(d)-(f) show the monostatic imaging results of the EPS block for 0° , 3° , and 6° rotations. Similarly to the metal target, when the EPS is oriented towards the radar in Fig. 2.39(d) is when the maximum received power is recorded. As the EPS is rotated, peak power is reduced. On the other hand, contrary to Fig. 2.39(a)-(c), the diffuse scattering component of the EPS produces very similar images in all three rotation cases, and the target does not decrease its average power down below the radar's detection threshold. Thus, no vanishing is experienced for this target, which even reveals its pellet-based composition through the circular glints in the images.

Although the metallic plate and EPS slab are simple targets, similar scattering behavior and vanishing effects can be observed in more complex items. This is specially critical when the target is a true threat, as is the case of a handgun. To illustrate this consideration a 9 mm

¹⁸Some powerful glints at ~ -45 dBm can be noticed throughout the sheet, likely due to cracks and scratches in the otherwise smooth surface.

Figure 2.39: Monostatic images of a smooth metallic sheet (a) oriented towards the TRX, (b) rotated 3° , and (c) rotated 6° . Monostatic images of a EPS block (d) oriented towards the TRX, (e) rotated 3° , and (f) rotated 6° . Colorbar: Returned power in dBm.

handgun, shown in Fig. 2.40(a), has also been scanned using the rotational positioner. In the initial unrotated monostatic image in Fig. 2.40(b), the power is distributed throughout the entire metallic parts of the handgun, such as the barrel and trigger. Meanwhile, the handle is rugged plastic, which creates diffuse scattering with lower peaks than the rest of the gun. In contrast, when the pistol is rotated [in Figs. 2.40(c)] the more powerful returns are present in the metal near the handle, which is now the only part of the gun oriented towards the radar. In consequence the barrel (which is now oriented away from the radar) is starting to vanish below the detection threshold, and the maximum returned power has decreased ~ 5 dB. Contrary to the smooth metallic plate, the barrel does not completely vanish due to its multi-faceted surface, which creates slightly diffuse reflection in addition to a powerful specular component. Very similar scattering effects have been observed by other authors at 620 GHz with a revolver [162, 163].

Figure 2.40: (a) Measurement target consisting of a replica Colt 1911 handgun held with a rotational positioner. (b) Imaging result when the target is oriented towards the radar. (c) Imaging result when the target is rotated just 3° . (d) Bistatic image of the handgun when radiated with a TX with $\theta = 6^\circ$ of angular diversity. (e) Multistatic image, combined from the previous monostatic and bistatic results. Colorbar: Received power in dBm.

These results pose a critical problem for a security imaging radar tasked with concealed threat detection. A subject hiding a gun or another weapon could avoid detection through

simple re-orientation of the item to shift it away from the radar's effective FoV. A solution to this problem is to increase the radar's effective FoV. This can of course be done by increasing the maximum FoV (via improvements to the reflector system or scanning mechanism), but also by modifying scene illumination. Because the effective FoV depends on the orientation of the target with regard to the transmitted waves, changes in the incidence angles of those waves will modify the effective FoV. In the rotated scenario in Fig. 2.40(c), if the gun is simultaneously radiated with an independent TX with $\theta = 6^\circ$ of angular diversity, the resulting bistatic image can be seen in Fig. 2.40(d). The power is once again distributed through the metallic areas and with peaks around -20 dBm, more alike Fig. 2.40(b) than Fig. 2.40(c). On the other hand, because the TX-to-target path does not employ the focusing reflector system, the cross-range resolution compared to Fig. 2.40(b) is deteriorated. In a fully multi-static implementation, this could be partially overcome with the combination of images from multiple viewpoints, as done in Fig. 2.40(e), from the results in Fig. 2.40(b), (c) and (d). The multistatic result presents worse resolution than the monostatic image in Fig. 2.40(b) (specially noticeable around the trigger, handle, and hammer), but better than the bistatic image in Fig. 2.40(d).

It can be seen from Fig. 2.40(e) that bistatic images lack the same quality as their monostatic counterparts, which can be related to their different spot sizes. The spot size in monostatic and bistatic configurations has been analyzed using a matrix of metallic spheres with a 2.5 mm radius, separated 7.5 cm in both directions, as shown in Fig. 2.41(a). A monostatic image of the target in Fig. 2.41(b) reveals detection and separation of individual spheres, with consistent received power levels. Meanwhile, the bistatic images in Fig. 2.41(c) and Fig. 2.41(d) present more deformed spheres. Notice that the monostatic case in Fig. 2.41(a) also presents some deformation, which could be due to the wooden picks holding the spheres and the plastic matrix in the background [106]. Comparing Fig. 2.41(c) and Fig. 2.41(d), this effect appears to worsen as the bistatic angle increases, with spheres reflecting less power and some starting to vanish.

Figure 2.41: (a) Matrix of 2.5-mm-radius metallic spheres. (b) Monostatic image. (c) Bistatic image of $\theta = 3^\circ$ angular diversity. (d) Bistatic image of $\theta = 6^\circ$ angular diversity. Colorbar: Received power in dBm. The total range of the colormaps is 9 dB in all images, but not set to common limits to account for different transmit power in bistatic measurements.

To offer a quantified metric, the 6 dB spot size in all three images in Fig. 2.41 has been measured. A 4x4 submatrix of spheres from each image (framed with a magenta rectangle) is used, and the mean width in the X and Y axis from their maximum returned power to its bilateral 6 dB drop-off is recorded. The same submatrix position has been chosen for all images, corresponding to one of the areas in the FoV where less deformation occurs. Table 2.11 summarizes the registered sizes, where the monostatic case has resulted in sizes equivalent to a previous experiment with the radar [102]. The results indicate similar sizes

among the bistatic scenarios, but specially deformed in the horizontal dimension. This is likely a side-effect of generating the bistatic images by introducing angular diversity only in θ , which rotates around the vertical axis.

Table 2.11: Mean spot size for the framed spheres in Fig. 2.41

Configuration	Y size (cm)	X size (cm)
Monostatic	1.87	1.56
Bistatic ($\theta = 3^\circ$)	1.87	1.66
Bistatic ($\theta = 6^\circ$)	1.88	1.68

Although these deformations can occur, in some cases, multiple monostatic and bistatic images can be combined to form a single image with higher quality than the individual contributions. This has been tested in a more realistic imaging scenario considering a nude female mannequin (representing a human target) shown in Fig. 2.42. The monostatic result is compared with a multistatic image combining monostatic and bistatic data. For the bistatic images, the TX has been placed at four locations around the radar to achieve $\theta = \pm 3^\circ$ and $\theta = \pm 6^\circ$ angular diversities. As can be seen from the figures, although this complex target is not as heavily affected by specular vanishing as the metal sheet or gun, the addition of the bistatic images fills some empty areas in the mannequin's surface where specular reflection has diverted returned power away from the RX. This effect is specially noticeable near the mannequin's chest, face, and lower stomach. Accumulative imaging results from the intermediate combinations are shown in Fig. 2.43. Increased visibility due to wider angular imaging, for a mannequin's head at 75 GHz - 82 GHz, has also been observed by other authors [168].

Figure 2.42: (a) Nude female mannequin. (b) Monostatic image when the mannequin is facing the reflector system. (c) Multistatic combination of all monostatic and bistatic images when the TX is placed at four locations achieving $\theta = \pm 3^\circ$ and $\theta = \pm 6^\circ$ bistatic angles. Colorbar: Range in meters.

Before discussing security imaging scenarios, the effect on point cloud quality of combining additional viewpoints to a monostatic image must be further studied. The analysis is performed on the point cloud data resulting from the scenario in Fig. 2.42. To quantify the improvement in quality obtained, a reference point cloud of the same nude mannequin is used as ground truth. This point cloud is generated with a structured-light scanner¹⁹ [169], which

¹⁹Device which emits a known IR pattern over the scene and estimates the reflective surface from the pattern's deformation. The model used is the Structure Sensor from Occipital, connected to a tablet.

Figure 2.43: Intermediate combinations from the imaging scenario in Fig. 2.42. (a) Monostatic only. (b) Monostatic + Bistatic at $\theta = -3^\circ$. (c) Monostatic + Bistatic at $\theta = \pm 3^\circ$. (d) Monostatic + Bistatic at $\theta = \pm 3^\circ$ and at $\theta = -6^\circ$. (e) Monostatic + Bistatic at $\theta = \pm 3^\circ$ and $\theta = \pm 6^\circ$. Colorbar: Range in meters.

allows mm-level spatial resolution. These types of depth sensors have been previously used as reference data in other multiview radars [218], and are also used in modern imagers for passenger control and tracking [92]. The reference point cloud can be seen in Fig. 2.44(a), where its higher quality when compared to radar-generated images is evident.

Figure 2.44: (a) Structured-light point cloud of a nude female mannequin. Colorbar: Range in meters. (b) Registered radar and structured-light point clouds of a nude female mannequin.

The median distance of each coordinate in the radar point clouds to its nearest structured-light neighbor has been used as the quality metric. The median is used to mitigate the contributions of outlier noise that would increase the mean distance. Assuming r_n as the n -th spot in a 3-D radar point cloud with N spots, having cartesian coordinates (x_n, y_n, z_n) , its distance to each point $[s_i, \text{having coordinates } (x_i, y_i, z_i)]$ in the structured-light cloud (M total points) is simply

$$D_{ni} = \sqrt{(x_n - x_i)^2 + (y_n - y_i)^2 + (z_n - z_i)^2}, \quad (\forall i | 1 \leq i \leq M) \quad (2.15)$$

where each radar spot's nearest structured-light neighbor is selected with

$$D_n = \min(D_{ni}), (\forall i | 1 \leq i \leq M) \quad (2.16)$$

and the final quality metric is

$$D_{med} = \text{median}(D_n), (\forall n | 1 \leq n \leq N) \quad (2.17)$$

As the radar point cloud grows denser when combining multiple viewpoints, and its missing areas are filled, the distance to the nearest structured-light neighbor will be decreased. In an ideal case, a perfectly-combined radar point cloud with mm-level cross-range resolution would produce near-zero distances to the reference data.

In order to accurately compute the distance between radar and structured-light point clouds without including any offsets, the two sets of data must first be registered. This has been manually done with the iterative closest point (ICP) algorithm [170], using radar data

from Fig. 2.42(b). The registered point clouds are shown in Fig. 2.44(b), where it can be seen that the ICP stage has achieved good coordinate association.

To objectively compare the image quality of combinations with a different number of viewpoints, the following sequence is used: I_1 is the single monostatic image from Fig. 2.42(b), I_2 is the combination of I_1 and bistatic images at $\theta = \pm 3^\circ$ angular diversity, and I_3 is the combination of I_2 and bistatic images at $\theta = \pm 6^\circ$ angular diversity. Table 2.12 summarizes the improvement from multistatic combination, where N_p is the number of points in the point cloud, D_{med} is the median distance between a radar coordinate and its closest structured-light neighbor, and the improvement in I_i is measured as the percentage reduction in D_{med} compared to I_1 . As can be seen from the results, adding the first two bistatic images at $\theta = \pm 3^\circ$ produces a considerable 19.5% reduction in D_{med} , while the next two images (at $\theta = \pm 6^\circ$) only increase that reduction to 20.7%. It must be taken into account that these results are only applicable to the mannequin, which presents a significant diffuse scattering component; for targets with a larger specular scattering component, successive images at increasing angular diversity should provide steady reductions in D_{med} .

Table 2.12: Effect of multistatic combination on image quality

Image	N_p	D_{med} (mm)	Improvement (%)
I_1	5965	8.2	-
I_2	17867	6.6	19.5
I_3	29804	6.5	20.7

In future analyses, additional quality metrics should be used to evaluate the improvement from multiview combination. To this end, specific tools to evaluate active imager quality have already been presented in the scientific literature [171]. Additionally, no-reference and reduced-reference metrics [172], common in photographic quality assessment, could also be evaluated.

2.3.2 Imaging Simulations

To further illustrate the effects of scattering and rotations in mm-wave images, two simulation tools have been applied to several imaging scenarios. These simulators allow fast, flexible, and inexpensive evaluation of complex imaging set-ups. This ability is specially valuable when dealing with equipment with high costs and lengthy development times, such as a 300 GHz radar.

The first simulator is a ray-tracing framework, based on the bidirectional reflectance distribution function (BRDF)²⁰, able to simulate the radar's scan [22]. This tool allows reproducing the imaging results of the 300 GHz radar in both monostatic and bistatic scenarios. The simulations are based on a ray-tracing approach to approximate the scan pattern, along with the use of a BRDF to model the scattering behavior of the target. Although the simulator does not use traditional electromagnetic propagation and interaction equations, ray-tracing combined with appropriate scattering functions has been previously proven to be an effective approximation [173]. A recent review of ray-tracing THz simulators can be found in [22]. The developed simulator uses previously-measured data from the radar, as well as a commercial depth sensor, to simulate new imaging scenarios. In published results, simulated scattering with this tool has shown very good agreement with actual measured data, for synthetic targets (such as a mannequin) as well as for live humans.

²⁰Which is a computer graphics concept, commonly used for surface rendering, which relates the reflected radiance of a surface to its incident irradiance.

Although ray-tracing might not be able to offer a result as accurate as Kirchoff or Maxwell computations²¹, it provides its results much faster. This advantage comes from the large electrical size of targets compared to the transmitted wavelengths, which forces electromagnetic simulators to take into account very fine surface meshes during simulation. This creates a computational bottleneck which can expand simulation time to hours or even days [173], contrasting with ray-tracing, which can be done in seconds or minutes. In addition, our simulator's code is implemented in a GPU, which further reduces the imaging simulation of a ~ 29000 -facet surface from 126.8 s in a serial execution to just 9.7 s in the parallelized version.

The simulation procedure for a given scenario is the following:

- First, the imaging scenario is defined. This includes the mechanical scan setting, the placement of the target, the reflector system, and all TRXs and TXs present in the scene, among other parameters. To do this, a virtual 3-D model of the target is required. The model must have sufficient quality (in terms of the size of its facets or voxels) in order for the ray-tracing approximation to provide accurate results. A minimum boundary for this quality would be the radar's cm-level resolution in range and cross-range.
- In the following step the simulator performs ray-tracing between all TXs, TRX's, and the target, for each pixel in the radar's scan pattern. This provides a list of which of the target's facets are visible to each TX or TRX, using the Möller-Trumbore ray-triangle intersection algorithm [175], which leads to computation of all the bistatic and monostatic paths present. Multipath reflections are not considered in this simulator.
- With the propagation paths set, the returned power for each ray is then computed. The returned power (P_r) is calculated with

$$P_r = \frac{P_t A_{eq}}{4\pi R^2} \cdot \sigma^o(\theta_i) \quad (2.18)$$

where P_t is the transmitted power, A_{eq} is the equivalent area of the receiving antenna, R is the distance to the target, and σ^o is the differential RCS [176], which is dependent on the incidence angle (θ_i). All parameters in the equation are pre-set during the scenario definition stage, except the differential RCS, which can be approximated by [177]

$$\sigma^o(\theta_i) = 4\pi f(\theta_i) \cos^2(\theta_i) \quad (2.19)$$

where $f(\theta_i)$ is a function modeled with the BRDF approach. To take into account both diffuse and specular scattering for any target, $f(\theta_i)$ is modeled by

$$f(\theta_i) = k \cdot f_d(\theta_i) + (1 - k) \cdot f_s(\theta_i), \quad 0 \leq k \leq 1 \quad (2.20)$$

where $f_d(\theta_i)$ and $f_s(\theta_i)$ are the diffuse and specular terms of scattering, respectively. To model diffuse scattering, Blinn-Phong's formulation is used [178], while for specular scattering, the Cook-Torrance model is selected [179]. Thus, the final approximation for the returned power in each pixel presents multiple model parameters (such as roughness distribution of the surface, weights for each type of scattering, and others), which must be manually set using previous imaging results from the radar. Once the returned power for each ray is obtained, an image can be formed by simply filtering out rays with power lower than the radar's detection threshold. In [22] some comparisons between actual images from the radar and simulated scenarios can be seen, highlighting the similarity between both sets of data.

²¹Broadly speaking, ray-tracing is a tool from the optical side of the spectrum, and accurate mm-wave simulations also require taking into account wave physics.

The simulated results for the same metallic sheet from Fig. 2.38 can be seen in Fig. 2.45. The sheet has been given a slight tilt backwards to match the measured image in Fig. 2.39(a), which makes the reflected power be centered somewhat upwards. In the rotated scenario in Fig. 2.45(b) the sheet's reflected power has dropped ~ 80 dB, much below the radar's threshold. Also, notice that because the simulation is optical in nature, border effects are not taken into account, and thus, the rotated simulation registers uniform power throughout the entire sheet.

Figure 2.45: (a) Monostatic ray-tracing simulation result when a metallic sheet is oriented towards the radar. (b) Monostatic ray-tracing simulation result when a metallic sheet is rotated just 3° . Colorbar: Received power in dBm.

A similar vanishing effect due to rotation in monostatic images can be seen in Fig. 2.46, where a 9 mm handgun has been simulated to compare with the measured results in Fig. 2.40. In the unrotated case in Fig. 2.46(b), the gun reflects constant and high power in most of its flat surface, with some angled facets redirecting less radiation and thus registering slightly less power. On the other hand, the rotated scenario in Fig. 2.46(c) presents much lower levels in the gun's flat areas (mostly barrel and trigger), while higher power is reflected from parts of the handle (which is now oriented towards the RX). Notice also that, in contrast with the sheet, the reflected power from the rotated barrel is still within the radar's threshold levels, due to a rougher surface, with more diffuse scattering than the sheet.

Figure 2.46: (a) 3-D model of a MP-443 handgun used as simulation target. (b) Monostatic ray-tracing simulation result when the handgun is oriented towards the radar. (c) Monostatic ray-tracing simulation result when the handgun is rotated just 3° . Colorbar: Received power in dBm.

The previous scattering scenarios have also been contrasted with simulations using the second simulation approach, a Physical Optics (PO) tool previously developed by the *Universidad de Vigo* [164, 165]. This custom-made software provides information related to electromagnetic interactions within the reflector system and electric field scattering from the target, which the ray-tracing simulator does not take into account. The PO code simulates the specific radiation pattern of the TRX through the reflector system, as well as the independent TX's COTS horn antenna. After that, the simulator models free-space propagation effects, as well as incident electric field and induced currents in the target, to finally predict their reflections towards the RX.

The first scenario tested with the PO simulator are the rotations seen in Fig. 2.38(c). A comparison between simulated and measured results is presented in Fig. 2.47, with 61 dB simulated losses in returned power after a 3° rotation for the monostatic case, showing good agreement with measurements.

Figure 2.47: Comparison between simulated and measured returned power from a smooth metallic sheet in monostatic and bistatic (at $\theta = 6^\circ$ and $\theta = 12^\circ$ angular diversities) placements, as the sheet is being rotated around its vertical axis. Measured data comes from Fig. 2.38.

Figure 2.48: (a) Simulated monostatic imaging result equivalent to Fig. 2.40(b). Simulated monostatic imaging result equivalent to Fig. 2.40(c) (a 3° rotation). (c) Simulated bistatic imaging result equivalent to Fig. 2.40(d) (a 3° rotation). Colorbar: Normalized received power in dB.

The 9 mm gun has also been simulated using the PO software, as seen in Fig. 2.48. The simulated images present the same scattering behavior as the measured results (strong

reflections from the smooth barrel when not rotated, and from the handle when rotated), validating the need for angular diversity when specular targets are rotated.

The scenario in Fig. 2.42, where the multistatic combination of images from a nude female mannequin was shown to present better quality than a single monostatic image, has also been simulated using the PO software, with the results seen in Fig. 2.49. Simulations once again show good agreement with the measured results, and the same vanishing effect around the stomach area from Fig. 2.42 can be seen in Fig. 2.49. As with the measured results, the combination of images from multiple viewpoints fills the missing detections.

Figure 2.49: (a) Simulated monostatic image equivalent to Fig. 2.42(b). (c) Simulated multistatic combination equivalent to Fig. 2.42(c). Colorbar: Range in meters.

Although custom-made software have been used in this section, several commercial alternatives²² could be used as simulators. Nonetheless, no single commercial product is currently able to accurately simulate all aspects of a multistatic imaging radar, which includes diverse elements, such as mechanical scanning, quasi-optical focusing, RF interactions, and base-band processing, among others. This is the main reason why authors using commercial software often resort to either combining various products, or modifying them to suit their needs [174, 180–182].

2.3.3 Evaluation of Multistatic Imaging for Security Screening

Specific security screening scenarios expanding on the generic results in the two previous subsections have been tested. For this application, the scenarios involve weapon concealment under different types of clothing. The first scenario is the female mannequin wearing a T-shirt [as shown in Fig. 2.50(a)], and hiding a ceramic knife taped to the abdomen [as seen in Fig. 2.50(b)]. The monostatic image when the mannequin is oriented towards the reflector system is presented in Fig. 2.50(c), and the monostatic image when the mannequin is rotated 10° counterclockwise is shown in Fig. 2.50(d). As can be seen from the results, when the knife's blade is oriented towards the reflector system the weapon can be detected as a protruding anomaly in the abdominal area. In contrast, the highly-specular nature of the blade (similar to the metallic sheet analyzed in Fig. 2.38) creates an empty zone in the

²²Such as Open FX, Zemax, FEKO, HFSS, or CST.

rotated image where the knife is placed, due to the lack of enough returned power to surpass the radar's threshold.

On the other hand, when the mannequin rotated 10° is simultaneously illuminated with a TX placed with $\theta = 20^\circ$ angular diversity, and its point cloud is combined with Fig. 2.50(d), the combined bistatic image [seen in Fig. 2.50(e)] retains the concealed threat. Although the knife in the bistatic image does not contain as many detections as the monostatic case, a blade-like anomaly can still be seen in the abdomen.

The following scenario deals with the same mannequin concealing the replica handgun from Fig. 2.40(a) in the abdomen area, as shown in Fig. 2.51(a) and (b). The monostatic image with the mannequin in normal orientation towards the reflector system can be seen in Fig. 2.51(c), where the weapon is clearly separated from the body. However, when the mannequin is turned just 10° counterclockwise, the gun in the monostatic image is not as clearly identified, as seen in Fig. 2.51(d). The handle, which presents a strong diffuse factor, is still visible, but the barrel has mostly disappeared due to specular scattering. On the other hand, when using the bistatic results from a TX placed with $\theta = 20^\circ$ angular diversity and the monostatic result in Fig. 2.51(d), the combined image in Fig. 2.51(e) still retains the gun just as well as the unrotated monostatic image.

Because the mannequin presents unrealistic scattering and penetration compared to a person, the final imaging scenario has been a test on a male human subject. In addition, to provide even more resemblance to an actual security screening case, the threat is concealed with a wool sweater instead of a light cotton T-shirt. Fig. 2.52(a) shows the scanned subject, who is hiding the replica handgun from Fig. 2.40(a) in the abdominal area. The monostatic image when the subject is oriented towards the reflector system is presented in Fig. 2.52(b), and the monostatic image when the human rotates 10° counterclockwise is shown in Fig. 2.52(c). Similarly to the knife scenario in Fig. 2.50, when the gun's barrel is oriented towards the reflector system the weapon can be detected as a protrusion from the subject's otherwise flat abdomen. When rotated, the barrel disappears from the image, although in contrast to the barrel and the knife, the gun's rugged handle remains detectable (as seen in Fig. 2.40). Meanwhile if the subject rotated 10° is simultaneously illuminated with the TX placed with $\theta = 15^\circ$, $\theta = 20^\circ$, and $\theta = 25^\circ$ angular diversities, the combined multistatic image from these placements and the monostatic image in Fig. 2.52(c) maintains the full weapon again, as seen in Fig. 2.52(d). This scenario has required more bistatic angles than Fig. 2.51, due to the added difficulties from the subject's flat torso and the thick sweater. To accurately combine all the resulting images, the bistatic results at $\theta = 15^\circ$ and $\theta = 25^\circ$ are artificially rotated to the $\theta = 20^\circ$ viewpoint in offline processing.

Figure 2.50: (a) Female mannequin wearing a T-shirt to hide a ceramic knife near the abdomen. (b) Placement of the knife underneath the T-shirt. (c) Monostatic image when the mannequin is facing the reflector system. (d) Monostatic image when the mannequin is rotated 10° counterclockwise. (e) Combined image of the mannequin from the bistatic result when radiated with a TX with $\theta = 20^\circ$ of angular diversity and the monostatic result in (d). Colorbar: Range in meters.

Figure 2.51: (a) Female mannequin concealing the gun from Fig. 2.40 with a T-shirt. (b) Photograph of the gun placement near the abdomen. (c) Monostatic image when the mannequin is facing the reflector system. (d) Monostatic image when the mannequin is rotated 10° counterclockwise. (e) Combined image of the mannequin from the bistatic result when radiated with a TX with $\theta = 20^\circ$ of angular diversity and the monostatic result in (d). Colorbar: Range in meters.

Figure 2.52: (a) Male human wearing a sweater to hide the handgun from Fig. 2.40(a) near the abdomen. (b) Monostatic image when the human is facing the reflector system. (c) Monostatic image when the human is rotated 10° counterclockwise. (d) Combined multistatic image of the male human from the bistatic results when radiated with TXs with $\theta = 15^\circ$, $\theta = 20^\circ$, and $\theta = 25^\circ$ angular diversities and the monostatic result in (c). Colorbar: Range in meters.

2.4 Future Work

This section summarizes some of the future and ongoing work related to improving the 300 GHz radar regarding real-time and multistatic imaging. First off, optimizations and future research lines for the real-time image processing architecture in Section 2.2 are presented. Lastly, some improvements to the multistatic imaging and scattering configuration in Section 2.3 are described.

2.4.1 Real-Time Imaging

The developed real-time processing system presents multiple opportunities for optimization, regarding aspects such as image refresh rate, cost, or user friendliness. Of all impacted areas refresh speed is still the top priority, since, if it is fast enough, is what allows accurate representation of the target's movements. As mentioned in the previous section, in our radar this speed is still limited by mechanical obstacles to approximately 2 fps. On the other hand, the imaging software has proven to be able to process scenes at a much higher rate (at least 9.35 fps), even allowing improvements up to around 46 fps (Fig. 2.26). Nonetheless, additional research lines to enhance this functionality can be examined.

2.4.1.1 GPU-based 3-D Reconstruction

As mentioned in the discussion of both algorithm versions, the surface reconstruction of the imaged scene is a bottleneck in the VTK code. This effect can be seen in Fig. 2.20 and Fig. 2.25, where the histograms for point-cloud-only representation register lower processing times by several orders of magnitude. The cause for this bottleneck is the large computational requirements of the Delaunay triangulation in 3-D. Although this reconstruction is currently implemented in serial VTK code, its algorithm can be parallelized into the GPU [153, 183]. The motivation behind this is the GPU's increased performance against CPU-only processing when an algorithm presents some type of parallelization opportunities. This is the case in

the Delaunay triangulation, where, for instance, the distance between points in the cloud, when assigning neighbors, can be computed independently.

Full integration of this GPU stage into the developed software has not been completed during this thesis, but some offline trials have been performed with a GPU-based Delaunay 3-D reconstruction implemented in open-source code called *gStar4D* [184], which reports speedup of up to five times when compared to the fastest 3-D serial triangulations. The code has been tested by feeding the NVIDIA GTX 970 GPU with two alternating datasets (each corresponding to a single point cloud image from our radar) until each dataset has been triangulated 1000 times. The first image (a point cloud of 2403 3-D coordinates) corresponds to a nude female mannequin, while the second (a point cloud of 3198 3-D coordinates) is from the same mannequin wearing a T-shirt over an improvised explosive device (IED). The performance of the code is timed internally without the use of external profilers, and the results may be seen in Fig. 2.53. The registered times (> 2.5 s) for this triangulation are much higher than the entire VTK code (< 0.5 s), as shown in Fig. 2.20 and Fig. 2.25, which means that the raw *gStar4D* implementation cannot be directly included in the processing algorithm without reducing image refresh. This can be due to the fact that in comparison to the usual point clouds in computer vision applications (which are commonly very structured), our radar point clouds present more noise and no clear structure, which can increase the difficulty when assigning neighbors and performing the triangulations. Nonetheless, this code can serve as a starting point for future parallelized implementations and further optimizations to adjust the algorithm could provide increased speed. Additionally, alternative reconstruction algorithms (not based on Delaunay triangulations) can be analyzed to select the fastest method.

Figure 2.53: Histograms of the registered processing times over 1000 iterations for the 3-D Delaunay triangulation of two radar images, when performed on the GTX 970 using *gStar4D*.

2.4.1.2 Multi-GPU Processing Architecture

A second, less critical improvement comes from the addition of another GPU into the PC. Including a second GPU into the hardware can reduce the workload imposed on the TITAN X device, by offloading into it part of the data in need of processing. This adds another layer of parallelism in the algorithm through concurrent processing of independent sections of the same point cloud in two different devices. The performance of the algorithm is then improved by reducing the total time required to process the cloud.

The main GPU manufacturers (NVIDIA and AMD) both support these multi-GPU architectures through the use of proprietary technology: Scalable Link Interface (SLI) [185], and

CrossFire [186]. These technologies allow multiple GPUs to communicate with each other, and the CPU, in order to parallelize an application's demands. Apart from the additional GPU, more hardware is required to implement the multi-GPU architecture, in the form of GPU-connecting bridges through which the GPUs communicate. These bridges, shown in Fig. 2.54, are necessary in order for the CPU to receive the processed data from the GPUs. This is because multi-GPU rigs usually work on a master-slave scheme, where one GPU is in charge of reporting back all processed information, even when both devices have treated equal workloads. This mechanism simplifies the CPU-to-GPU, and viceversa, communications. The bridge allows the slave GPU to send its processed data to the master, which in turn delivers it back to the host PC.

Figure 2.54: Photographs of four interconnected NVIDIA TITAN GPUs (a) and two AMD Radeon GPUs (b). The NVIDIA models are connected with an SLI bridge, while the AMD devices are connected with a CrossFire bridge.

Duplicating GPU devices is not usually done in simple parallel computing architectures, because it increases the complexity of data traffic and kernel scheduling, while providing a marginal improvement in performance²³. More often than not, if the initial GPU is not fully resource-limited during processing, a thorough optimization of a single GPU device can yield the desired enhancement. Meanwhile, multi-GPU architectures require deep knowledge of asynchronous memory transfers and locking mechanisms, among others. Additionally, in a multi-model multi-GPU rig, the more powerful GPU will be limited to the speed and memory of the slowest model (in order to simplify communications) [188]. Furthermore, not all GPU models are compatible with the SLI and CrossFire technologies. Due to PCIe slot constraints (discussed in Annex C), a multi-GPU trial has not been performed. Additionally, in our case, using the GTX 970 and TITAN X jointly would limit the latter to the processing capabilities of the former, which does not guarantee the desired improvement. Nonetheless, in the future, if two GPUs of the same model are available, a multi-GPU architecture could provide some slight improvement to algorithm speed. More substantial improvements could come from the recent updates to multi-GPU technologies that both manufacturers are undertaking. NVIDIA has begun replacement of SLI with a new product called NVLink [189] (which reportedly improves communication speed and ease), while AMD has moved onto bridge-less (communicating simply over PCIe) architectures with DirectX 12's API-based support and abstraction.

²³Rarely does duplicating the number of GPU devices double the performance, and the improvement diminishes as more devices are added [187].

2.4.1.3 Hardware Optimizations

Along with the software-based future work mentioned, multiple RF and mechanical optimizations can be performed on the system. For instance, before further increasing the software's image refresh speed above the mechanical movement's limit, enhancing that movement's speed should be the priority. As discussed in Section 2.2.3, the current scanning mechanism based on high-speed mirror movement is limited to approximately 2 fps. To solve this issue, most recent imagers used combinations of mechanical movement with multi-TRX RF chains (reducing the scanning time per image by a factor equal to the total number of TRXs) [92, 110, 135].

Incrementing the number of TRXs will substantially elevate system cost. In our radar, some previous modifications to the mechanical set-up can still be performed to speed-up the scanning without adding another RF chain. These can be summarized as follows:

- **Mirror balancing.** Currently, the scanning mirror's weight is not perfectly balanced around its rotation axis. This leads to very noticeable, and possibly hazardous, vibrations when speeds above 8000 rpm are demanded from the rotation motor, which are desired if the scan is to be performed faster. To solve this, the mirror must be mechanically adjusted to compensate its weight imbalance.
- **Removal of reduction drive.** The union between the rotation motor and the scanning mirror is done via a 1:4 reduction drive²⁴, which allows incrementing the torque in the mirror (at the expense of lowering the rotation speed to one fourth of the motor's). This means that, for example, when the rotation motor moves at 8000 rpm, the mirror is only spinning at 2000 rpm. This device is included to speed-up the time required for the mirror to reach scanning velocity, which is longer without the drive. By removing this drive, the mirror's speed will be increased by a factor of four, at the cost of reaching its scan speed during start-up at a slower pace.
- **Reduction of chirp time.** To ensure enough pixel density when increasing mirror speed, chirp time must be reduced for the radar to spend less time illuminating a single spot. The current chirp time is 150 μ s, but the DDS used has been fully tested to lower times with the full band sweep. Further analysis needs to be performed in order to study the effects of reducing chirp length on other system areas, such as transmitted power or phase noise.
- **Updated version of elevation motor.** At its fastest configuration, the elevation motor used is current-limited to 1.2 A, corresponding to a mirror-tilting speed of 1800 steps/s with a linear movement of 0.0317 mm/step. To increase this speed, which is the main factor limiting FoV scan time, a newer version of the elevation motor must be installed. For example, a different version of the same motor model presents a linear movement of 0.127 mm/step.

With these considerations in mind, a comparison between the scan patterns generated at different scan speeds may be seen in Fig. 2.55. In said figure, a high-quality image scan (at 1.17 s per scan) is compared against the current fastest scan (at 0.55 s per scan) and a scan achievable with the proposed modifications (at 0.13 s per scan). The scan with the optimizations presents a pixel count (1791 spots) that is half of the current fastest scan's (3583 spots) and less than one fourth of the high-quality's (7800 spots), which will noticeably reduce point cloud quality. This can be seen in the figure's bottom row, where the imaging results of a nude female mannequin for each scan pattern are shown. As expected, the decrease in point cloud density is noticeable, but the target can still be seen. To avoid any

²⁴Model PLE40 from Neugart.

decrease in image quality, the reduction in density could be overcome with the use of the filtering or surface reconstruction techniques presented throughout the thesis. Additionally, at the speed achieved (7.7 fps, for 0.13 s per scan) subject movement would be represented more accurately. It is worth mentioning that additional common mechanical limitations, such as motor overheating from continuous use, would also need to be addressed.

Figure 2.55: FoV scan patterns for a (a) high-quality image (8000 rpm, 827 steps/s, 150 μ s per chirp), (b) the fastest scan currently available (8000 rpm, 1600 steps/s, 150 μ s per chirp), (c) a modified maximum (16000 rpm, 1800 steps/s, 75 μ s per chirp), (d)-(f) and their corresponding imaging results for a nude female mannequin.

2.4.2 Multistatic Scattering and Imaging

The main focus of future work regarding multistatic imaging architectures should be defining the optimal placement of the transmitters in order to provide detections under all threat orientations with the lowest cost possible [174]. Since the cost of RF equipment in the mm-wave and submm-wave bands remains high, reducing the number of TXs required is a priority in order to keep the architecture cost-effective. For this task, both simulators used in Section 2.3.2 would be useful tools. The PO simulator will allow to predict scattering behavior in a given scenario at an electromagnetic level, while the ray-tracing model could provide fast estimations of visibility conditions for the targets present in the scenario. Additional approaches to determine the most efficient multistatic architecture using, for instance, convex optimization models, have been presented [190].

Aside from this development, several other minor improvements could be made in relation

to the proposed evaluation:

- A simple expansion to the experiments presented in this section would be the addition of new multistatic viewpoints. In the images scenarios discussed, angular diversity was achieved by revolving the TX around only the target's vertical axis, and at fixed height and range. Assuming a spherical coordinate system centered on the target, the scenarios presented have only introduced diversity in the azimuthal angle. Nonetheless, the radial distance and polar angle can also be varied with the tripod's adjustable height, pitch, yaw, and roll. Although roll is not expected to influence the results (due to the TX's horn symmetry), variations in pitch, yaw, height, and range, would allow adding complementary angular diversity.
- Although RF hardware costs at 300 GHz are still a restriction, research in MMIC technology is rapidly decreasing this limit. As detailed in the next chapter of this thesis, several COTS radar front-ends at mm-wave frequencies have been recently available at fractions of the cost of a single waveguide component. The ability to place full front-ends in low-cost chips not only allows acquiring a larger number of TXs in the proposed multistatic architecture, but also implementing newer configurations, such as the walkthrough system described in Fig. 2.31(b).
- Additional work should be performed to expand both simulator's abilities. In the case of the current version of the PO simulator, the main limitation is the simulation of only perfectly-conductive materials (although the *Universidad de Vigo* is currently working on solving this issue). Bypassing this obstacle is necessary to simulate dielectric threats, and take into account the existence of penetrable range layers in the scene. For the ray-tracing simulator, one of its main restrictions is the lack of multipath propagation, which is abundant in multistatic scenarios.
- In a future true multistatic configuration²⁵, registration of the multiple point clouds generated might not be as straightforward. Similarly, if point cloud combination is expected to be performed online as the radar scans, or different scan patterns are simultaneously used, the registration process will require redesign. For all of these scenarios, the use of the structured-light depth sensor could be expanded to serve as a real-time scene reference, as has already been proposed by other authors [155].
- Keeping in mind the good agreement between simulated scenarios and their measured counterparts, the simulation tools used in this section could also be utilized to develop a database of synthetic imaging results. The benefit of using the simulators instead of measured data for this task would be the speed, ease, and flexibility of the software tools in comparison to building an actual imaging set-up for all scenario variations considered. Since simulated images show nearly-identical results to measured scenarios, the database could be used, for example, to train automatic detection algorithms (as previously proposed in [22] and explained in a later chapter of this thesis).
- In the presented evaluation, an unexploited aspect of a multi-TX or multi-TRX architecture is the use of multiple-input and multiple-output (MIMO) imaging concepts [191–193]. The use of orthogonal waveforms and/or multiplexing techniques would allow better separation of the individual propagation paths than with simple range windowing, while avoiding interferences. Approaches such as sparse topologies could be used to reduce the number of receivers in such a configuration [155].

²⁵Where multiple TXs are simultaneously radiating and not simply virtualized with different placements of the same RF equipment.

Chapter 3

122 GHz and 94 GHz Multi-sensor Security Imager

This chapter is focused on the implementation of a hybrid system using active sensors (in the form of 122 GHz FMCW radars) along with passive sensors (94 GHz passive RXs) and an optical depth-sensing device. As with the 300 GHz radar in the previous chapter, the hybrid imager's application is also security screening. The simultaneous use of multi-type sensors acts as a solution to the individual imaging obstacles of each individual sensor type, as will be demonstrated with various screening scenarios. All sensors used are based on COTS devices and the reflector system used has been modified from a pre-existing product.

The sections in this chapter are organized as follows:

- In first place, Section 3.1 presents the MMIC around which the 122 GHz radar has been developed. Following this introduction, the radar's characterization in several performance areas (such as resolution, bandwidth available, or stability) is included.
- After introducing the active sensor, Section 3.2 details the motives behind the development of the hybrid imager. This section also includes descriptions and characterizations of the remaining elements of the imager, such as the reflector system, the passive sensors, and the optical device. An image fusion algorithm to combine the images from each sensor is also demonstrated in this section. Lastly, several imaging scenarios showcasing the hybrid system's imaging performance are presented.
- The last part of the chapter, in Section 3.3, lists some of the future research lines and ongoing work related to the developments previously presented.

3.1 122 GHz Radar Characterization and Performance

3.1.1 122 GHz MMIC-based Radar

One of the main challenges in evolving a high-frequency imaging radar, such as the one described in Chapter 2, into a real-world commercial product is the reduction of the system's volume. In mature technologies, larger devices usually demand higher manufacturing and handling costs, as well as increased power consumption. Up to recently, this has been the case for high-frequency RF systems, where bulky components have been historically available, but integrated devices were at first limited to academic research only. However, with the rise of

MMIC technology, multiple commercial solutions compacting entire radars into a single chip have become accessible with low prices. The integration of the full RF front end (from the VCO to the antennas) in a chip allows a considerable volume reduction in comparison to traditional waveguide-based solutions (such as the 300 GHz system in Chapter 2). Although the initial design and prototyping costs for these chips are high, scale economies as the technology matured have allowed manufacturing thousands of MMICS into single wafers with similar performance among all chips and minimal need for manual tuning and assembly [15], greatly reducing the fabrication costs. Price-reductions in the front end provide a large effect in the overall cost of a system through the decrease of RF cost, which is usually one of the most expensive areas.

In accordance to current cost- and size-reduction trends in mm-wave systems [86], a 122 GHz radar has been developed from a commercial MMIC front end. The motivation behind the development of this low-cost and integrated radar is twofold: to assemble an inexpensive front end to make radar imaging arrays commercially-feasible (while maintaining overall performance), and to produce a portable sensor to be used in alternative applications, such as NDT, vital sign monitoring, automotive radars, or sensor networks. The MMIC chosen is designed and manufactured by Silicon Radar GmbH, and has been acquired along with a complete FMCW radar kit (which includes the MMIC, a dielectric lens, a baseband board for signal conditioning, and a microcontroller for signal processing). Additional details regarding the commercial kit, and the modifications performed to adapt it into the final FMCW sensor used throughout this thesis may be found in Annex D. The commercial kit itself is shown in Fig. 3.1, and a comparison among other available products is shown in Table 3.1.

Figure 3.1: Photographs of the commercial FMCW radar kit from Silicon Radar GmbH (a) without and (b) with its dielectric lens attached. Modified from [194].

Several initial considerations can be made regarding the selected product:

- The decrease in operating frequency and bandwidth from the 300 GHz system is evident. Assuming similar antenna structure, such a decrease has considerable effect in both cross-range (through the larger wavelength) and range (through the narrower bandwidth) resolutions in imaging applications, which could ultimately cause an inability to accurately detect concealed threats. Although FMCW imaging systems have historically decreased their operating frequency from around 500 GHz - 600 GHz, these have usually settled around the 300 GHz band, which provides a good tradeoff between cross-range resolution and clothing transmissivity. On the other hand, several imaging systems in W-band frequencies (75 GHz - 110 GHz) and even lower bands have also been presented with demonstrated threat-detection capabilities.
- Nonetheless, the sensor's reduction in size and cost when compared to the 300 GHz RF chains and microwave source is also evident. In addition, Silicon Radar also distributes

Table 3.1: Comparison among commercial MMIC solutions integrating full TRXs

Manufacturer	Model	f_c (GHz)	BW (GHz)	P_{TX} (dBm)	Antennas	Size (mm ²)	$P_c^\diamond$ (mW)	Phase Noise* ($\frac{dBc}{Hz}$)
Silicon Radar [†] [194]	TRA_120_01	122	7	-3	Patches	8x8	380	-90
Silicon Radar [†] [194]	TRA_120_012A	122	16	-3	Dipoles	8x8	440	-90
Silicon Radar [194]	TRX_024_007	24	3.6	4	(Not included) [⊙]	3x3	89	-102
TI [195]	AWR1642	79	4	12.5	(Not included) [⊙]	10.4x10.4	1300 - 2140	-95 - -93
TI [195]	AWR1443	79	4	12	(Not included) [⊙]	10.4x10.4	1720 - 2100	-95 - -93
TI [195]	AWR1243	79	4	12	(Not included) [⊙]	10.4x10.4	1620 - 2010	-95 - -93
Analog Devices [196]	EVAL-RADAR	24	0.25	8	(Not included) [⊙]	(N/A)*	>340	-108
Infineon [197]	BGT70	73	5	12	(Not included) [⊙]	6x6	350 - 480	-80 [‡]
Infineon [197]	BGT80	83	5	12	(Not included) [⊙]	6x6	50 - 480	-80 [‡]
ST [198]	STRADA431	24	0.25	13	(Not included) [⊙]	6x6	270	-75 [‡]
Fraunhofer [199]	(N/A)	93	15	5	(Not included) [⊙]	(N/A)*	4500	-100

[⊙]Power consumption.

*At 1 MHz offset.

[†]Acquired for this thesis.[⊙]Not included directly on chip.

*Not a single chip (evaluation board or full sensor).

[‡]At 100 KHz offset.

its MMICs separately from the commercial radar kit, allowing even further system integration and cost savings. Although an initial compaction of part of the 300 GHz RX chain (up to ~ 50 GHz) via MMIC technology had been designed [127] (details in Annex G), this update would still require some bulky waveguide components. In contrast, Silicon Radar’s MMIC integrates a full RF front end (including VCO, multipliers, amplifiers, mixers, and antenna) in a 8×8 mm² chip costing ~ 150 €.

- Because Silicon Radar’s MMICs are derived from an European research project ended on mid-2013 [200], the commercial availability of the devices is fairly recent. At the time of acquisition, this led to a relative lack of product documentation (aside from some research papers and concise datasheets) and third-party evaluations by widespread clients. Namely, official measurements regarding lens beamwidths, IQ imbalances, noise figure, or phase noise along the RF chain were missing. For the developments in this thesis, this absence of data forced a dependence on reverse-engineering, along with private communications with Silicon Radar. During this thesis, the first MMICs were acquired in 2016, and future acquisitions continued up to late-2018. It must be noted that in this span, official and publicly-available documentation from Silicon Radar regarding its products considerably improved (including updated datasheets, user manuals, and new measurements).

The final architecture of the in-house modified 122 GHz radar and a photograph detailing its components may be found in Fig. 3.2. The radar has been mainly modified by removing Silicon Radar’s microcontroller board and designing a PCB to house a new microcontroller and some signal conditioning stages. The modifications enable continuous FMCW transmission, real-time sampling and processing, sampling locked and triggered with signal generation, and farther range detections; all functions necessary for the imaging application, and previously missing in the commercial kit. Modifications are described in Annex D.

Figure 3.2: (a) Block diagram of the architecture of the modified 122 GHz radar from a Silicon Radar GmbH commercial kit, and (b) photograph detailing its components.

The modified radar is composed of three stacked boards: the top board houses the MMIC, while the bottom two include various RF components and electronics, such as filters, amplifiers, and a microcontroller. The microcontroller is tasked with waveform configuration and communication with a PC, which performs all signal sampling and processing. The chirp

signal is generated with a phase-locked loop (PLL) built around a Phase-Frequency Detector (PFD) and a VCO at 60 GHz. The VCO, antennas, and I/Q mixers are all found in the MMIC itself, while baseband components (such as the PFD, its reference oscillator, and several filters and amplifiers) are located in the two bottom boards. A two-antenna configuration is used (one for transmission and the other for reception), as is common in FCMW systems, to provide adequate isolation between the transmit and receive channels. A sample of the PLL's reference oscillator is used as the ADC's sampling clock for coherence, while an end-of-ramp pulse from the PFD is used to trigger the ADC's sampling and thus synchronize it with chirp generation.

At the time of writing this thesis, Silicon Radar GmbH currently sells multiple MMIC products [194], with varying characteristics (such as center frequency, bandwidth, or number of channels). The modified radar allows the use of different MMICs from Silicon Radar GmbH, as long as they share a similar pinout. Because of this, during this thesis two different MMICs (shown in Fig. 3.3) have been used: TRA_120.01 and TRA_120.012¹, whose corresponding datasheets can be found online². TRA_120.01 is a front end centered at 122 GHz with an advertised bandwidth of up to 7 GHz, while TRA_120.012 shares the same center frequency with a bandwidth of up to 16 GHz. Their RF architectures are the same as shown in Fig. 3.2, albeit with different VCOs and antennas³. A more detailed comparison between the two chips can be seen in Table 3.2.

Figure 3.3: Photographs of the (a) TRA_120.01 and (b) TRA_120.012 MMICs from Silicon Radar GmbH, reproduced from [194].

Table 3.2: Comparison among the two Silicon Radar GmbH MMICs used in this thesis

Parameter	TRA_120.01	TRA_120.012
Center frequency (GHz)	122	122
Available bandwidth (GHz)	7	16
Antenna type	Patch	Dipole
Power consumption (mW)	380	440
Phase noise ($\frac{dBc}{Hz}$)*	-90	-90
Transmitter power (dBm)	-3	-3
Receiver Gain (dB)	8	8

*At a 1 MHz offset.

Both chips are meant to be used at the 122.5 GHz Industrial, Scientific, and Medical (ISM) radio band, whose licensing depends on local governance. As of 2017 in Spain, where

¹Version A, out of four possible variations from A to D.

²Regularly updated at <https://siliconradar.com/products/>

³TRA_120.01 uses two grids of on-chip patch antennas, while TRA_120.012 uses two on-chip dipole antennas.

this thesis has been developed, explicit gubernamental authorization is needed to operate in this band, as transmissions may interfere with other devices [201]. In addition, the ISM band at 122.5 GHz only has 1 GHz of allocated bandwidth (between 122 GHz and 123 GHz), which means that when transmitting at full bandwidth, both MMICs exceed the regulation limits. Since the measurements and tests performed during this thesis have been performed in laboratories and other controlled environments, the full capabilities of the MMICs have been used without interfering with other ISM users. For a future commercial use of the applications developed in this Section, specific licensing arrangements must be carried out.

3.1.2 Evaluation of Radar Performance

In imaging radars, three main undesired effects limit the performance of a system:

- **RF non-linearities** [202]. Non-flat amplitude and phase responses throughout the components in the RF chain create amplitude and phase modulations in the transmitted and received signals. When computing the spectrum, this effect produces an increase in the side-lobe levels near the target's beat frequency. If the increase is high enough, less powerful targets can be masked, reducing the radar's dynamic range.
- **Transmitted phase noise** [203]. If the transmitted phase noise is more powerful than the thermal noise, the noise floor in the RX during transmission will rise, reducing the SNR of a target. Because the phase noise is related to the transmitted power, increasing the chirp's power will not improve SNR⁴. Phase noise will once again be specially noticeable near the target's beat frequency, whose main lobe will be severely widened.
- **Spurs**. Spurious signals usually occur during the signal generation stage, and are multiplied, amplified, and transmitted along with the chirp. In the received spectrum, they can be misinterpreted as targets, which enforces the use of range windowing to isolate the true beat frequency. In PLL-based generation, such as with Silicon Radar's MMIC, these signals are less common, due to the operation of the PLL's feedback loop, which progressively compensates VCO's spurs within its bandwidth.

To evaluate the radar's performance, several figures of merit can be defined:

- **SNR**. A measurement of the ratio between an ideal target's returned power and the noise floor, which can be dominated by transmitted phase noise or RX thermal noise. In an imaging radar, the measurement scenario usually involves a small isolated metallic sphere, providing a powerful point-like target. A similar parameter is the instantaneous dynamic range (IDR), which is a measurement of the radar's ability to simultaneously detect targets with different power levels. This figure of merit will of course be affected by the noise floor, but also the side-lobe levels of the most powerful reflection, which can mask the spectral peak of the secondary target.
- **Resolution**. A measurement of the radar's ability to separate equally-powerful targets in all dimensions. In an imaging radar, two different resolutions can be defined, depending on their orientation. In the range axis, the resolution is determined by (1.20), via the transmitted bandwidth. In the two remaining axes, the resolution depends on the operating frequency and the antenna's dimensions.

⁴This is in contrast to when the system is thermal-noise-limited, where increased power creates higher SNR.

- **Ranging Accuracy and Sensor Stability.** A measurement of the radar's overall precision when performing range measurements. This can be defined as the deviation between the ideal distance associated to the detected beat frequency and the target's actual range. This precision is deteriorated by the radar's instability in several areas, such as transmitted power, phase, and frequency. To thoroughly determine the accuracy these instabilities must also be measured.

3.1.2.1 SNR

For a preliminary analysis of the capabilities of the developed sensor, simulations of the received power and the SNR at the RX's output have been performed following (1.23) and (1.24). Both the traditional radar equation and its imaging variant have been used due to the sensor's capability to include a dielectric lens during measurements, which limits the MMIC's radiated beamwidth in a similar fashion to that of an imaging reflector system. Because the transmitted beamwidth is not taken into account in both equations, the simulations are valid for both MMICs (assuming equal values in all other parameters). The RX bandwidth (B) in both cases is set to the width of the FFT bins [which is related to the transmitted chirp time through (1.19)]. A sweep of possible transmitted power (P_t) values is performed from -7 dBm to 1 dBm, according to official datasheets. The transmit (G_t) and receive (G_r) gains are approximated to 23 dB, including the lens, with measured data from [204]. Meanwhile both the noise figure (F) and antenna temperature (T_a) are set to typical values of 5 dB and 290 K, respectively. Finally, the propagation losses (α) are approximately $1 \frac{dB}{km}$, according to the attenuation windows presented in [205]. The resulting graphs may be seen in Fig 3.4 for the SNR using both equations. Only short ranges (from 0.1 m to 10 m) have been considered in order to limit the representation within a typical range window for imaging systems. In both cases the simulated target is a metallic corner reflector with a peak RCS of 70 dBsm, and the dwell time is equal to a single 150 μs chirp. As expected the received power for the analysis according to (1.24) is lower than according to (1.23), due to the elimination of the effects of G_t and σ . However, for an imaging application with the sensor, following (1.24), SNRs between 50 dB and 60 dB can still be expected. This is a critical benchmark of the system, as SNRs of around 40 dB have been proven necessary for high-quality images in previous FMCW imaging radars [97].

Figure 3.4: SNR analysis using the MMIC-based sensor at 122 GHz for a 70 dBsm corner reflector, according to (a) (1.23) and (b) (1.24).

To test out the simulations, the SNR of the radar (with the TRA_120.012 MMIC) has

then been evaluated in a measurement scenario using a real 70 dBsm corner reflector. The measurement has been made over a small range of distances, suitable for the intended security application. A 150 μ s chirp with a bandwidth of 15.1 GHz is transmitted for each range point. The evolution of the SNR as the metallic corner reflector is placed at increasing distances may be seen in Fig. 3.5, where for each distance the target's reflections are captured and processed with a windowed FFT to represent the spectrum. As is expected from a conventional radar system following (1.23), received power decreases as the target moves farther away with a $\frac{1}{R^4}$ dependence on the range to the target⁵.

Figure 3.5: (a) Scaling of the SNR of a metallic corner reflector with a peak RCS of 70 dBsm as it is placed at increasing distances from the radar. Sweep time is 150 μ s, and 15.1 GHz of bandwidth are transmitted. (b) Photograph of the corner reflector used.

However, the measured SNR values and the ones seen in the simulations differ by large margins. This is mainly due to the powerful leakage from the TX into the RX, which manifests itself as a powerful phase noise curve slowly descending from dc to higher frequencies. This phase noise raises the noise floor above the RX's thermal noise, estimated by the radar equation. The radar's noise floor has been measured by aiming at a mm-wave absorber at approximately 1 m, and has been compared against the noise floor of the external ADC used for signal sampling, loaded with a 50 Ω impedance. The comparison is presented in Fig 3.6, where the TRA_120.012 MMIC has been used with a bandwidth of 15 GHz and varying sweep times. When comparing the two cases, one of the first considerations that can be made is the lack of flatness in the MMIC's spectrum, specially at lower frequencies, which will clearly reduce the dynamic range when measuring targets at close distances. The curve is common in FMCW radar systems, where leakage from the transmit antenna interferes in the receiving antenna [206], creating a powerful and constant masking at near-zero frequency. Another noticeable feature are the spurious signals detected at various frequencies, such as 4 MHz, 5 MHz, and 8 MHz. Because the corresponding spectral peaks are not present in the ADC noise floor, the spurs must come from the radar's components. The effect on the noise floor of modifying the chirp time can also be seen in the figures, which is not constant because an FFT with a bandwidth inverse to the chirp time is used to generate the spectrums. As the chirp time is increased, the FFT's bandwidth decreases, reducing the amount of noise allowed into the spectral representation.

Additional minor misestimations in some radar equation parameters, specifically the noise figure and antenna temperature, which have been set without measurement data, could have also influenced the large difference between SNR simulations in Fig 3.4 and measurements in

⁵Minor disagreements can be due to reflector orientation.

Figure 3.6: Noise floor measurements when the external ADC is connected to (a) a 50 Ω load and (b) the developed radar aiming at a millimeter-wave absorber. Varying chirp times are considered and used during the FFTs that produce the represented spectrums.

Fig. 3.5. In any case, the high RCS and large dimensions of the corner reflector make it an unusual target in actual imaging scenarios. To offer a more realistic approach, the SNR when measuring typical materials in security screening (such as fabrics, explosives, or human skin) has also been simulated. In addition, the SNRs are also calculated for varying dwell times, to analyze the effect of modifying the chirp time. The materials are listed in Table 3.3, and the results from the simulations may be seen in Fig 3.7, where each material's RCS (σ) is proportional to its reflectivity (r) through

$$\sigma \propto 20 \log_{10}(r) \quad (3.1)$$

Table 3.3: Materials commonly present in security imaging scenarios

Material	Reflectivity*
Explosive	0.24
Metal	1
Skin	0.35
Denim	0.01
T-shirt	0

* Approximated from measurements at 100 GHz in [20].

To analyze the radar's sensitivity, a single measurement of a point-like target (metallic sphere with a 2 cm diameter) is compared against the noise floors from Fig 3.6. The sphere has been placed at approximately 1 m and the resulting comparison is presented in Fig 3.8. As can be seen from the figure, when reflections from a target are present, the noise floor is further elevated, likely due to its phase noise contribution. The target's peak (at approximately 1.2 MHz) still rises 33 dB above noise, although its IDR is reduced to 25 dB due to the side-lobe levels of the target. If the target's peak is compared against the noise floor of the MMIC when no target is present in the scene (labeled $SNR_{phasenoise}$ in the figure) the SNR would increase to 53 dB. In addition, if the phase noise from TX leakage could be reduced to the ADC floor, the SNR would be further increased to 77 dB. In an updated version of the radar (outside of the main scope of this thesis), both the leakage phase noise and the target phase

Figure 3.7: SNR analysis using the MMIC-based sensor at 122 GHz for a collection of common imaging objects, according to (1.23), when the dwell time is (a) $75 \mu\text{s}$, (b) $150 \mu\text{s}$, and (c) $1500 \mu\text{s}$.

(c)

Figure 3.8: (a) Sensitivity measurement and noise floor comparisons, (b) and zoom, for a metallic sphere with a 2 cm diameter illuminated by the 122 GHz radar with a 6.3 GHz bandwidth and a $150 \mu\text{s}$ chirp. (c) Photograph of the measurement set-up.

noise have been reduced⁶ through the improved performance of a new reference oscillator for the PLL.

Of course, the SNR analysis presented is only valid when the radar is working in far-field region conditions for its radiation. The lower boundary for this area is known as the Fraunhofer distance [207]. For a range r to be considered part of the far-field region, it must confirm

$$r > \frac{2D^2}{\lambda} \quad (3.2)$$

where D is the aperture diameter of the radar's antenna, and λ its signal's wavelength. In the case of the MMICs considered, the physical dimensions of the antennas are below 1 cm. However, the use of the focusing lens (whose diameter is around 1.5 cm) modifies the effective aperture of the sensor. As illustration, the Fraunhofer distance for the lower and upper frequencies delimiting a 18 GHz bandwidth around the MMIC-based radar's center frequency is presented in Fig 3.9, for aperture diameters up to 5 cm. For the diameters considered, the Fraunhofer distance lies below 2.5 m at all frequencies, which is an acceptable limit for short-range imaging applications.

Figure 3.9: Fraunhofer distance according to frequencies delimiting a 18 GHz bandwidth around the developed radar's 122 GHz center frequency.

3.1.2.2 Resolution

High-frequency FMCW imaging systems usually have diffraction-limited spatial resolution [108], meaning the size of the smallest features representable in an image is related to the wavelength of the transmitted signal. This resolution can be defined as the angle (θ) at which the first minimum of the Airy disk is found [208], which is calculated as

$$\theta = 1.22 \frac{\lambda}{D} \quad (3.3)$$

where λ is the transmitted signal's wavelength, and D the antenna diameter. On the other hand, from a purely trigonometric perspective, the small-angle approximation for a sinusoid allows estimating the resolution by

⁶At least 40 dB less for offset frequencies up to 10 MHz.

$$\theta \simeq \frac{d}{R} \quad (3.4)$$

where d is the radiated spot size at a given range R . Taking into account both equations, the radar's spot size at any distance can be formulated as

$$d \simeq 1.22 \frac{\lambda R}{D} \quad (3.5)$$

With this in mind, the MMIC-based radar's spot size at 122 GHz has been estimated for distances ranging from 0 m to 10 m, and antenna diameters from 1 mm to 10 mm (corresponding to the order of magnitude of the MMIC's physical dimensions). The resulting graph can be seen in Fig 3.10(a), which shows that even for low ranges, the spot size can still be in the order of meters. This is undesirable for a security imaging application, where the finest detail is needed when representing an image, to confirm whether or not a threat is being concealed. For this reason a focusing lens with a larger aperture than the antennas, such as the one provided with the commercial kit from Silicon Radar GmbH, should be used during measurements. The addition of the lens narrows the radiated beamwidth, providing a smaller spot size, as can be seen in the estimation in Fig 3.10(b) (where lens diameters from 5 mm to 50 mm have been used). The data show that, when using the lens, for ranges below 4 m the estimated spot size is below 1 m for all but two lens diameters.

Figure 3.10: Spot size estimations for the MMIC-based sensor at 122 GHz according to (3.5), with (a) and without (b) using the focusing lens.

Meanwhile, in the range dimension, the radar's resolution is determined by its maximum transmitted bandwidth, through (1.20). To measure this parameter, the radar is configured for full-bandwidth transmission, and a high-frequency probe is placed in the divided signal path of the feedback part in the PLL loop (as shown in Fig 3.11). The bandwidth cannot be measured at the VCO's output due to lack of access inside the MMICs. Nonetheless, the divided signal can be multiplied by a factor of 64 to provide a representation of the bandwidth transmitted at the RF band centered around 122 GHz. This measurement has been performed for both of the available MMICs, with the results seen in Fig 3.12. As can be seen, the maximum available bandwidths are 6.3 GHz for TRA_120_01 (from 119.2 GHz to 125.5 GHz) and 15.1 GHz for TRA_120_012 (from 115.7 GHz to 130.8 GHz), which are slightly lower than the reported 7 GHz and 16 GHz bandwidths, respectively, in official datasheets. Additionally, it can be seen that the bandwidth for TRA_120_01 suffers from a progressive

attenuation from its lowest frequency to its highest, with a maximum loss of 3.2 dB at 125.5 GHz. This effect is likely due to the lower performance of the VCO or frequency divider as frequencies increase, which seem to have been improved or equalized for the TRA_120_012 (a more recent MMIC product).

Figure 3.11: Transmitted bandwidth measurement by placing a high-frequency probe in the feedback loop of the front end's PLL.

Figure 3.12: Maximum available bandwidth measurement for the (a) TRA_120.01 and (b) TRA_120.012 MMICs. To perform the measurement a probe is placed at the divided signal path of the PLL loop, and the frequency axis is converted to the RF range.

To measure the actual range resolution of the MMIC front ends, a simple evaluation scenario is designed using two rectangular sheets, one made from cardboard and the other from aluminium. These are placed facing the radar, with the cardboard in front of the metal. Because the cardboard is easily penetrated by mm-wave radiation, two target peaks are expected in the captured spectrum. By correctly orientating the metal sheet, the two peaks can then be made equally powerful, creating a scenario with two equipotential targets at different ranges. The distance between the two items is then lowered to its minimum value that still generates a null of at least 6 dB between their spectral peaks. This distance is noted as the minimum range necessary to separate two equally powerful targets, and is also known as the range resolution [205]. The results for both MMICs are found in Fig 3.13, where each MMIC is configured to transmit the maximum bandwidths measured in Fig 3.12. The TRA_120.01 MMIC registers at least a 4.91 cm resolution, while that value goes down to 2.33 cm for the TRA_120.012 chip. Both of these magnitudes are slightly larger than the 4.31 cm and 1.80 cm, respectively, theoretical resolutions according to (1.20). Nonetheless,

it can be seen (specially for the TRA_120_012 MMIC) that some range and power margins are still available in the null between peaks. This means that the targets can still be brought a little close together, fully reaching the theoretical performance.

Figure 3.13: Range resolution measurement for the (a) TRA_120_01 and (b) TRA_120_012 MMICs. To perform the measurement two targets (one cardboard and one metal) are placed facing the radar, one in front of the other, and their spectrums captured. The resolution is defined as the minimum distance that still ensures a 6 dB null between targets in the spectrum.

3.1.2.3 Ranging Accuracy and Sensor Stability

From (1.18) the relationship between the beat frequency and a target's range is determined. Through this equation, all detected peaks in the radar's spectrum are translated into a measured distance. However, the equation only accounts for a theoretical and ideal performance. To quantify the possible deviations from this relationship, and correct them in future experiments if necessary, a sequence of measurements is performed. In it, the 70 dBsm corner reflector is placed at increasing distances, while the beat frequency generated is registered. At the same time, the theoretical range corresponding to that frequency is calculated with (1.18), and compared against the measured data. The results for both MMICs, while transmitting a 6.3 GHz bandwidth with 150 μ s chirps, are presented in Fig 3.14. As can be seen, only small and inconsistent deviations are present in the measured data, indicating that their most likely source are involuntary errors during the measurement set-up.

Having confirmed that no beat frequency offsets are present, a measurement sequence to quantify the residual range error at varying distances is performed. For this sequence, the 70 dBsm corner reflector is placed at increasing distances over a 2 m swath, and the measured range is compared against the actual distance to the radar. Results are shown in Fig 3.15 and averaged in Table 3.4. Chirp period is 150 μ s, while the transmitted bandwidth is 6.3 GHz for TRA_120_01 and 15.1 GHz for TRA_120_012. In the measurement sequence, several different methods for measuring the range have been considered, detailed below:

Table 3.4: Mean error in mm for both MMICs with all measurement methods in Fig 3.15

MMIC	Mean	Single	I	Q	Padding	Freq. + Phase
TRA_120_01	0.3655	0.3655	0.3725	0.3591	0.3647	0.3556
TRA_120_012	0.1131	0.1132	0.1142	0.1095	0.1124	0.1093

Figure 3.14: Measurement of the radar's deviation from (1.18) in its beat frequency to range relationship. A 70 dBsm corner reflector is placed at increasing distances and its beat frequency is registered and compared against the theoretical value from (1.18), for the (a) TRA_120.01 and (b) TRA_120.012 MMICs.

Figure 3.15: Measurement of the radar's ranging accuracy. A 70 dBsm corner reflector is placed at increasing distances and its measured range according to (1.18) is compared against the actual distance to the radar, for the (a) TRA_120.01 and (b) TRA_120.012 MMICs.

- **Mean of 10000 chirps.** In this scenario, for each range distance a 1.5 s data capture is performed. With 150 μs chirps, this corresponds to 10000 individual range measurements per distance. For each chirp, the range measurement is performed through a peak detection on the spectrum resulting from an FFT of the complex sum of the radar's I and Q channel data. These measurements are then averaged to form a single value per distance.
- **Single chirp.** In this scenario, a single 150 μs chirp is used to measure the range in each sequence step. When compared against the previous scenario, this case serves as illustration of the radar's range stability. If the averaged measurement differs largely from the single measurement, the radar has undesirable variation when measuring a static target.
- **I channel only.** In this scenario, instead of performing an FFT on the chirp's complex sum of the I and Q channels, only the I channel data are used. This case serves as study of the need of sampling both baseband channels to offer more accurate measurements.
- **Q channel only.** This scenario is analogous to the previous one, using the Q channel instead of the I.
- **Increased zero-padding.** In this scenario, zero-padding on the spectrum used to measure range in each step is increased by a factor of 100. This means to determine whether the increase in precision due to higher padding is critical for reduced error.
- **Frequency + Phase.** In this scenario, the phase information of the FFT is used along with its magnitude. The magnitude of the FFT allows separating the spectrum in ΔR range bins, within which the phase of the corresponding spectral peak can provide additional resolution [209] through

$$\Delta d = \frac{\Delta\phi \cdot c}{4\pi f_0} \quad (3.6)$$

where Δd is the distance corresponding to a $\Delta\phi$ phase difference, c is speed of light in vacuum, and f_0 is the central transmitted frequency. In contrast to the range resolution, this range variation is not dependent on the transmitted bandwidth, only the operating frequency and phase drift.

Several considerations can be made from the results on the figure. First, the error axis for the TRA_120_012 MMIC is far below the levels for the TRA_120_01 MMIC. This is of course due to the increased bandwidth of the TRA_120_012 chip, which provides more than double the range resolution. In addition, it can be seen that in both cases the use of the phase information for additional resolution achieves the better performance throughout the entire sequence. On the other hand, the use of only one of the two baseband channels (I or Q) provides different results. While using only the I channel results in the worst performance, using only the Q channel provides comparable results to the frequency plus phase method. This would indicate that there exists some form of I/Q imbalance in the radar, which degrades the measurements when simultaneously using both channels. Lastly, both the use of averaging of multiple chirps and an increased padding provide marginal benefits when compared to the single chirp case. This indicates the high range stability of the radar, and the lack of need for additional padding (which results in faster processing). The results in Fig 3.15 show similar performance to previously-published data of equivalent experiments done with Silicon Radar GmbH's commercial kit [210].

The measurement of the radar's accuracy when illuminating a single target can then be extended to a scenario where multiple targets are present in the scene. For this purpose, another sequence is performed, in this case with two equally-powerful targets placed in the

same line of sight, but at different ranges (such as in Fig 3.13). In this sequence, the first target is placed at 2 m, while the distance between targets is progressively increased. Meanwhile the radar measures this distance by detecting the reflections from both targets. In this experiment, the error is defined as the difference between the radar-measured distance from target to target, and the actual distance between them. The results for both MMICs are presented in Fig 3.16, which show better performance (~ 10 mm of maximum error) than a previously-published experiment with Silicon Radar GmbH's commercial kit (~ 30 mm of maximum error) [211].

Figure 3.16: Measurement of the radar's range-measuring accuracy when two targets in the same line of sight are present in the scene. The first target is placed at 2 m, while the second target progressively moves farther away. For the error calculation, the measured distance between targets is compared against the actual distance, for the (a) TRA_120.01 and (b) TRA_120.012 MMICs.

Another important area for radar stability is in its phase performance. As illustration, the Fig 3.17 shows the distance errors due to different theoretical phase fluctuations at increasing frequencies, according to (3.6).

Figure 3.17: Distance error for different phase fluctuations and frequencies, according to (3.6). For each frequency and fluctuation combination, target movements smaller than the corresponding curve cannot be measured accurately.

The phase stability of the radar can be defined as its phase drift over a given period of time for a static target. For a moving target, as stated in (3.6), a variation in phase for a reflected signal can be associated to the target's movement in the range direction. If the radar's phase stability is low, and thus the registered phase for any given measurement has large variation, this could forbid the accurate measurement of small target movements.

At high-frequencies, such as the 122 GHz band, the phase stability is affected by three main factors. First, the RF components themselves create phase modulations in the transmitted signal due to temperature fluctuations⁷. In addition, atmospheric variations in pressure, humidity, and temperature cause air permittivity changes, which include further phase shifts (although mostly present in outdoor scenarios). Lastly, the target's own vibrations and small-amplitude movements can cause a phase modulation on the returned signal.

To quantify the radar's phase performance, several measurements have been performed. In the first scenario, a metallic plate attached to a mechanical XY positioning stage with micrometer-level precision from Parker Hannifin Corporation, as shown in Fig 3.18(a). The initial measurement consists of simply registering the equivalent distance error associated to the phase variation according to (3.6), for the static plate during a 3 s capture. The results can be seen in Fig 3.18(b), where a 150 μ s chirp is used for each datapoint, and 6.3 GHz of bandwidth are transmitted with TRA_120_01.

Figure 3.18: (a) Photograph of the positioning stage and metallic plate. (b) Distance error from phase variation, according to (3.6) for the static metallic plate during a 3 second capture using 150 μ s chirps. (c) Histogram of the distance error.

⁷Specially notable in active components, and the system's reference clock, which is multiplied various times through the RF chain.

The maximum variation in distance is $9.29 \mu\text{m}$, which would indicate that, for the corresponding MMIC and waveform configuration, displacements smaller than that variation cannot be accurately measured. Notice how, although the registered fluctuation presents peak-to-peak variations, no downward or upward trend is detected, meaning that the phase is noisy but stable.

To confirm whether larger displacements are measurable, the XY positioning stage is used to move the metallic plate further back during a continuous capture with the radar. The stage's resolution, determined by a turn of its knobs, is $10 \mu\text{m}$, meaning that only movements larger than or equal to that distance can be performed. In Fig 3.19 a series of movements is measured with the radar, in the same MMIC and waveform configuration as Fig 3.18.

Figure 3.19: Movement sequence with the XY stage, as captured by the radar with $150 \mu\text{s}$ chirps. In (a) the plate is moved back $10 \mu\text{m}$, the minimum resolution of the stage. In (b) the plate is moved back another $50 \mu\text{m}$. In (c) the plate is moved 1.5 mm in approximately 17 smaller steps.

If larger chirp times are used, phase performance improves, as less noise is integrated in the narrower FFT bandwidth. To verify this, the static metallic plate is again radiated with the TRA_120_01 MMIC transmitting 6.3 GHz of bandwidth, but this time using $600 \mu\text{s}$ chirps. The resulting distance variation is shown in Fig 3.20, where the previous maximum distance variation from a static target ($9.52 \mu\text{m}$ in Fig 3.18) has been reduced to $3.52 \mu\text{m}$ ⁸.

Figure 3.20: (a) Distance error from phase variation, according to (3.6) for the static metallic plate during a 0.7 second capture using $600 \mu\text{s}$ chirps. (b) Histogram for the distance error.

⁸One fourth of the best mark obtained with the sensor in the literature [204]).

Notice also that, compared to Fig 3.18(c), the standard deviation of the distance error has been reduced by approximately half, as expected due to now using four times the chirp time (and dividing by four the FFT bandwidth).

Aside from the short-term measurements presented, during continuous operation in the radar is capable of maintaining its phase performance. This is proven by the long-term results of Fig 3.21, where the radar is left continuously transmitting during approximately 1 hour, and phase measurements of the static metallic plate are taken every 5 minutes. As can be seen from the graph, although the mean phase values slightly rise over time (a 0.053 rad phase drift in 60 minutes), their maximum variation within each measurement step stays relatively similar. In this sequence, the largest registered intra-capture variation is 0.019 rad, which would correspond to $3.72 \mu\text{m}$ according to (3.6).

Figure 3.21: Distance error from phase variation, according to (3.6) for the static metallic plate during 1 s captures, taken every 5 minutes, using $600 \mu\text{s}$ chirps.

However, it has been noticed that during continuous operation the radar's baseband and RF components warm up, as shown by the IR images in Fig 3.22. Because temperature fluctuations within the front end can be cause for phase variation in the transmitted and received signals, a simple experiment where the radar's temperature is decreased has been performed.

Figure 3.22: IR images of the 122 GHz radar's PCB, taken every 5 minutes, as it is in continuous operation. In a span of 15 minutes the radar's maximum temperature increases by 2°C . As can be seen, the USB connection of the microcontroller (which feeds the power consumption of the entire radar) and the MMIC are the two hottest components. Colorbar: Temperature in $^\circ\text{C}$.

To do so, three PC ventilation fans are used, aiming at the radar's PCB from one side,

as shown in Fig 3.23(a). No atmospheric compensation is done because the measurements are performed indoors, where little variations in pressure and humidity are expected. The comparison between using or not the fans during a phase measurement of the static metallic plate may be seen in Fig 3.23(b) and (c), for the TRA_120.01 MMIC transmitting 6.3 GHz of bandwidth using 600 μ s chirps. From these figures it may be seen that the use of a temperature stabilizing mechanism noticeably decreases the peak-to-peak distance error, which has been reduced from 2.59 μ m to approximately 1.91 μ m.

(a)

(b)

(c)

Figure 3.23: (a) Photograph of the temperature control set-up using PC ventilation fans. (b) Distance error from phase variation, according to (3.6) for the static metallic plate during 1 second captures using 600 μ s chirps, with and without temperature control. (c) Histograms of the distance errors with and without temperature control.

Since the RF components' efficiencies are expected to vary under temperature fluctuations, so will their output power. Therefore, because of the cascaded architecture of the front end (where a component's output is the next component's input), the final transmitted power is also expected to change. These changes in radiated power have been evaluated by means of the same temperature-control set-up shown in Fig 3.23. Due to the unavailability of an anechoic chamber for the traditional set-up used to measure transmitted power, the assessment has been based on returned power. Assuming a static target and waveform configuration, all returned power variations can be associated to the corresponding transmitted power changes. The results for a hour-long continuous capture of the reflections from a metallic plate, for both MMICs, can be seen in Fig 3.24. Although without temperature control the variations tend to stabilize after a warm-up period of a few minutes, they do so with considerable losses

(~ 2 dB for TRA_120.01 and ~ 5 dB for TRA_120.012, which is likely larger due to the effect of temperature fluctuations in its increased bandwidth). On the other hand, the controlled scenarios show very similar and low losses (below 0.3 dB).

With these results and the phase stability measurements in mind, it is recommended to use a temperature-control system in phase measurements. An effective solution is the use of a thermal chamber to electronically set and monitor ambient temperature. A more cost-efficient solution would be the use of a digital thermometer placed in the radar's PCB, connected in a feedback loop to ventilation fans in order to lower the temperature whenever it exceeds a threshold.

Figure 3.24: Comparison, with and without temperature control, of returned power variations from a static metallic plate using $600 \mu\text{s}$ chirps with 6.3 GHz of bandwidth for the TRA_120.01 MMIC and 15.1 GHz of bandwidth for the TRA_120.012 MMIC.

In addition to phase uncertainty and power reductions, fluctuations in temperature may cause amplitude and phase modulations in the transmitted and received signals. This is partially caused by the dependence of the efficiencies of RF components with regard to their temperature, which can generate non-flat power responses and temperature-dependent phase delays across the transmitted bandwidth. This has been an effect previously seen in other high-frequency FMCW radars [126], including the 300 GHz imaging radar presented in Section 2.1 [102]. If a received signal is affected by these modulations, the corresponding spectrum is widened and the main sinc's side lobes are increased, causing a loss in range resolution and instantaneous dynamic range. To compensate the effects of these modulations, a calibration procedure can be applied to a received signal [121]. This procedure consists of referencing the received spectrum to that of a point-like target⁹ at the same distance. Although the calibration procedure produces a substantial improvement in spectral purity, the reference spectrum must sometimes be refreshed to preserve a perfect compensation, as seen for the 300 GHz radar in Fig 3.25, specially during the radar's initial warm-up period. Note that, in the figure, after a 1 hour interval the calibration already produces a deteriorated compensation.

The effect of applying the calibration procedure to the developed 122 GHz radar, for both MMICs, can be seen in Fig 3.26. For the corresponding measurement, a metallic sphere with a 2 cm diameter is placed at 1 m and is self-calibrated. The non-calibrated and calibrated spectrums are then compared to analyze the procedure's effects, and an ideal sinc at the same frequency is included for further comparison. Observing the resulting spectrum, it is noticeable that the non-calibrated sinc does not seem to be affected by non-linearities,

⁹Usually an isolated small metallic sphere.

as it resembles very closely the reference sinc¹⁰. In addition, the calibrated sinc does not seem to improve the spectral benchmarks, further indicating that no modulations are to be compensated. The non-existence of these modulations is likely due to the high integration of the RF front end into a single chip, which produces an equally distributed temperature between all RF components, as well as improved adaptation and identical TX and RX paths.

Figure 3.25: Variations in the calibration procedure's performance in the 300 GHz imaging radar, when aiming at a static target, as time passes. The same reference signal is used for calibration in all times.

Figure 3.26: Spectral comparison of non-calibrated, calibrated, and ideal sincs from a metallic sphere placed at 1 m, for the (a) TRA_120_01 MMIC transmitting 6.3 GHz of bandwidth in 150 μ s chirps and the (b) TRA_120_012 MMIC transmitting 15.1 GHz of bandwidth in 150 μ s chirps.

3.1.3 122 GHz Waveforms Considered

Throughout the thesis, usually only three waveform configurations have been considered, in an effort to reduce the extension of the document. The configurations for these waveforms are as follows:

¹⁰With more reduced side-lobe levels, through the use of a windowing function.

- **Configuration 1.** This is the waveform resulting from the TRA_120.01 MMIC when transmitting 6.3 GHz of bandwidth with 150 μ s chirps.
- **Configuration 2.** This is the waveform resulting from the TRA_120.012 MMIC when transmitting 6.3 GHz of bandwidth with 150 μ s chirps.
- **Configuration 3.** This is the waveform resulting from the TRA_120.012 MMIC when transmitting 15.1 GHz of bandwidth with 150 μ s chirps.

3.2 Development of a Hybrid Imaging System

The 122 GHz radar developed in Section 3.1 has been used as the active RF source for a new security imager [105]. The imager is implemented as a hybrid system, combining active and passive mm-wave sensors, along with a depth-sensing camera used as a noise-removing tool for each sensor's images. Each type of sensor is able to compensate some of the imaging limitations of the other, thereby offering increased detection performance when compared to single-type imagers.

A large number of commercial single-type security imagers are passive systems, partly due to their simpler architecture. Nonetheless, passive systems are affected with some imaging limitations, such as temperature contrast in indoor environments and detector cooling for sensitivity, which are not present in active imagers. Although these factors are not critical in active imaging, active-only systems are also limited in some imaging functionalities, such as the specular reflection problems mentioned in Section 2.3. A straightforward solution to overcome individual weaknesses can be the combined use of both sensors. Temperature-related passive imaging is not dependent on transmission incidence angles, which allows it to provide image data when an active system might not. On the other hand, active imaging is indifferent to the contrasts of threat, body, and background temperatures, enabling results in cases where passive imaging fails. In a hybrid system, an FMCW radar can provide active 3-D image data, useful for separating clothing layers and shapes, to be fused with transmission-independent 2-D data from a passive sensor. A diagram of a sample hybrid set-up is shown in Fig. 3.27.

Figure 3.27: Diagram of simultaneous active, passive, and optical scanning of a human target.

Even though some modern hybrid imagers exist, as summarized in Table 3.5, these are still constrained in some aspects. In [212,213], no optical sensor is included and sensor fusion techniques are not applied (additionally, in [212], the FMCW radar is still not integrated for measurements). In [214] an image fusion process is applied, but without an optical image to

remove unwanted features and facilitate threat detection. In addition, recent funding from European institutions aims to create new and updated hybrid systems. In [91,92] the active and passive sensors of the CONSORTIS project have been tested individually, but have yet to fuse their information. The TERASCREEN project [93] is also expected to deliver a hybrid system, although no security imaging scenarios have been presented yet.

Table 3.5: Summary of current state-of-the-art high-frequency hybrid imaging systems

Reference	Active* (GHz)	Passive* (GHz)	Resolution [†] (cm)
Our work	115.7 - 130.8	94	~2.5
Robertson <i>et al.</i> [92]	325 - 355	250 & 500	1
Dill <i>et al.</i> [212]	24 - 25.5	28 - 36 & 84 - 96	≥14
Dolganova <i>et al.</i> [214]	250	250	1.2
Herschel <i>et al.</i> [93]	345 - 375	94 & 220 & 360	0.5 [‡]

* Active and passive frequencies of the systems.

[†] Cross-range for the active sensor.

[‡] Expected value, not yet measured.

Figure 3.28: (a) Representation of the hybrid imager's reflector system including the active, passive, and optical sensors. (b) Photograph of the side-by-side passive (gold) and active (red) sensors in front of the metallic mirror.

In the hybrid imager described in this Section, the developed 122 GHz radar acts as the active sensor, while a COTS 94 GHz RX acts as the passive detector, and a Microsoft Kinect camera serves as a depth-sensing support device. Some academic and commercial imagers have also included optical sensors as support devices for their single-type imagers. This addition is specially common in commercial passive imagers [119,215,216], where the optical device is mainly used for simple threat-visualization to avoid privacy concerns. In active systems, optical sensors have also been incorporated, for example to reduce the number of voxels during surface reconstructions by limiting the represented area [217], for passenger control functions [92], or to calibrate and attach multi-view scans [218]. In the our imager, the Kinect is used to remove noise and unnecessary information from the outputs of the other sensors, as well as facilitate threat visualization. Meanwhile, the active and passive sensors are placed side-by-side in an array distribution in a reflector system provided by a commercial imaging system, as shown in Fig. 3.28.

The sensors are combined at a hardware stage (being integrated in a single imager) as well as at software level (via simultaneous processing and sensor fusion techniques). The different processing chains for each sensor are executed in parallel, and their resulting images are combined using fusion techniques. Additionally, the use of commercial products provides a reduced-cost alternative to custom-made equipment. A summary of the hybrid imager's main parameters is presented in Table 3.6.

Table 3.6: System parameters for the developed hybrid imager

Parameter	Value
Standoff distance [†]	1.7 m
FoV	100x200 cm ²
Active frequency band	115.7 - 130.8 GHz
Passive frequency	94 GHz
Chirp period	150 μ s
Range resolution	1.8 cm
Temperature resolution	4°
Transmitted power	~3 dBm
Phase stability [‡]	0.018 rad
Minimum scan time	1 s
Processing time	< 1 min
Estimated cost	~60.000 €

[†]Measured from the imager's front panel.

[‡]Short-term stability (< 1 second), measured peak-to-peak.

3.2.1 Initial Imager and Passive Receiver Characterization

The reflector system used for the hybrid imager comes from Wavecomm's 94 GHz passive imager [219], shown in Fig. 3.29(a) and (b).

Figure 3.29: (a) Exterior and (b) internal photographs of Wavecomm's passive imager, (c) and diagram of the FoV scan achieved with the combined movements of the reflector system.

The reflector configuration used achieves a FoV of approximately 2 m vertically and 1 m horizontally, at a standoff distance of 1.7 m from its front panel. The imager uses a flat metallic mirror and a paraboloid reflector (both 70 cm in diameter) to redirect the radiation incoming from the target point to another point where the sensor is placed. The location of the target point and the sensor point are geometrically related by the parameters of the paraboloid (focal point, focal length and axis direction) and by the inclination of the mirror. The flat mirror has the effect of producing a virtual focus and virtual paraboloid axis as the specular image of the physical ones. The scanning of the FoV is then performed via two synchronized movements: a fast rotation of the mirror's axis produces a flapping motion which provides vertical scanning of the target point, while the RX platform's slow lateral movement allows for horizontal scanning of the target. The combined motion creates a raster scan of the FoV, as depicted in Fig. 3.29(c).

As shown in Fig. 3.30, a full turn of the mirror performs two vertical FoV sweeps in opposite directions: first the FoV is scanned downwards, and then upwards. Although the reflector system has a maximum speed of 1 fps when using a sensor array, unless otherwise stated the images here presented have a scanning time of 4 s.

Figure 3.30: Ray-tracing illustration showing the vertical FoV sweep performed in a single mirror turn in four steps in chronological order, from (a) to (d).

Image forming follows a different scheme for each sensor. For the active radar, the image is generated as the reflector system scans the FoV by mixing a sample of the transmitted signal with the round-trip-delayed reflection from the target in each pixel. The passive sensor also uses the scanning movement, but in this case the value in each pixel is simply the amplitude of the sampled signal. In the mirror's motor, an encoder's data is used to organize both sensors' pixels into a matrix, where each column is a turn of the rotating mirror. While the

mirror sweeps the FoV, the Microsoft Kinect photographs the scene in front, and performs depth sensing via a structured-light approach [220].

The passive RX is located on top of the moving platform, in front of the reflector system and beside the active sensor, as indicated in Fig. 3.28. The device is a 94 GHz PMMW-10-0001 imaging front-end from Farran Technology [221], with a W-band horn antenna from Flann Microwave Ltd [222]. The RX is supplied with 5 V at 42 mA and outputs a 1 V dc signal proportional to the temperature of the scanned object, which is amplified and conditioned in a custom-made board before being sampled with a dedicated ADC (although the active sensor's ADC could also be used). The sensor's output is confirmed when measuring a 50 Ω WR-10 load¹¹ at 25° attached to its RF port instead of an antenna. The recorded output may be seen in Fig. 3.31.

Figure 3.31: (a) Passive sensor's output when terminated with a 50 Ω load at 25°. (b) Histogram of the sensor's output.

A PCB is used to condition and sample the sensor's raw output, in a low-frequency integrated circuit (IC) design whose block diagram may be seen in Fig. 3.32.

Figure 3.32: Block diagram of the PCB used to condition and sample the passive sensor's output.

The sensor's output is directed to the PCB via a coaxial cable with an SMA connector. The PCB's first stage is an RC low-pass filter with a 7 KHz cutoff frequency designed to eliminate high-frequency noise from the image. The RX's integration time (τ) is listed as 1 μ s [221], but this value is only valid when capturing the RX's raw output without the use of our PCB. When using the PCB, the low-pass filter increases the integration time to 143 μ s (the inverse of its bandwidth). This change affects the sensor's sensitivity (ΔT) [221], which at a temperature of 25 °C (T_N) is calculated as

$$\Delta T = \frac{T_A + T_N}{\sqrt{\beta\tau}} \quad (3.7)$$

¹¹Metrology Bronze Termination 27045.

where T_A is 295 K and β is the RX's bandwidth (the full W-band; 35 GHz). Without the PCB ($\tau = 1 \mu\text{s}$) the sensitivity is listed as 4 K, while with the PCB ($\tau = 143 \mu\text{s}$) it would be just 0.26 K. After filtering, the next stage in the PCB is a differential RX needed to transform the RX's differential output into a single-ended signal. In addition, this stage raises the output's mean level from approximately 1 V to around 2 V to properly take advantage of the ADC's 0 V - 5 V range. Two amplifiers are then used to add gain to the output within the ADC's limits. Amplification factors are user-adaptable to calibrate each individual sensor's slight differences in mean output level. Finally, an ADC samples the filtered and amplified output at a frequency configured by the user. The digital samples are passed on to an Arduino microcontroller which in turn sends them to a PC via USB connection. Outputs for the PCB stages can be compared in Fig. 3.33, where a 0.9 V - 1.1 V sinusoid has been input into the PCB for amplifier calibration (simulating the RX's advertised maximum and minimum voltages). Notice how the differential RX raises the mean level, and how the amplifiers add gain (~ 12 dB total) to the signal without saturating the ADC's 5 V input limit.

Figure 3.33: Outputs at different stages of the PCB diagrammed in Fig. 3.32 when a 0.9 V - 1.1 V sinusoid is input for testing and calibration.

3.2.1.1 Temperature Resolution

The front-end's temperature resolution is advertised as 4 °C at 25 °C [221]. A measurement scenario to confirm this figure of merit has been created with the use of a frozen 0.5 L plastic water bottle and a handheld IR camera¹². The bottle is placed 10 cm away from the passive RX, with mm-wave absorber as background, while the IR camera is overlooking the scene. The bottle, starting at 0 °C, is left to naturally heat up to room temperature (22 °C), while RX and IR measurements are taken every 5 minutes. Results are shown in Fig. 3.34, where each datapoint averages a 1 second capture with an oscilloscope at 10 KS/s. The non-linear appearance of the data is explained by the exponential nature of heating processes and, partially, inaccurate targeting (the RX may be illuminating the absorber, surrounding air, or hotter objects through side lobes) and the heterogeneous temperature distribution throughout the bottle. Nonetheless, a polynomial fit to the data accurately shows a rising slope as the temperature increases. Temperature resolution can be calculated from

$$\sum_{i=j}^{j+n} (p_{i+1} - p_i) \geq \sigma_{max} \quad (3.8)$$

¹²Model E60 from FLIR Systems.

where p_i is the value of the polynomial fit in datapoint i , and σ_{max} the maximum standard deviation measured in the sequence (2.4 mV). This formula states that the minimum accurately-measured temperature variation must be larger than or equal to the maximum noise in a measurement. The resolution is then the variation in temperature of the n datapoints that imply a jump in voltage larger than σ_{max} . In this sequence, the maximum n (worst case scenario) is 15 datapoints at around 9 °C and the resulting resolution would be 7.5 °C¹³. If the resolution is instead measured with datapoints at 22 °C (in the more linear region of the polynomial, and also nearer to the manufacturer's temperature during measurement) the value decreases to 3.5 °C.

Figure 3.34: Sequence of passive RX output measurements over time, as the target (a water bottle initially at 0 °C) gradually heats up to room temperature (22 °C). The X-axis marks the minimum temperature registered in the bottle.

The sensor's ability to separate a human from a threat is illustrated in Fig. 3.35, where the RX has been used to image four different scenarios.

Figure 3.35: Histograms of the passive RX output for four different scenarios (a mm-wave absorber, a human hand, a room-temperature water bottle, and a clear sky). The temperature of each target is listed as reference.

The sensor's output is represented in an histogram approximating its probability density

¹³Because datapoints are measured every 0.5 °C.

function in four target scenarios: a mm-wave absorber (as a typical background reference), a human hand, a room-temperature water bottle, and a clear sky (as an ideal background). Each of the presented histograms is the result of a single second-long measurement while the target is static in a room-temperature empty space, to avoid noise due to movement, ambient temperature, or surrounding objects. As can be seen in the figure, the raw RX outputs of the first three scenarios present a clear separation in mean value, which must still be enhanced through later processing.

The sensor's outputs when aiming at different targets are further compared in Fig. 3.36. Because in the intended application the sensor must separate humans from background and concealed threats, the three targets considered have been a room-temperature mm-wave absorber (acting as homogenous background), a human hand, and a room-temperature water bottle. Each datapoint in the curve is averaged over a 1 s capture with the passive RX. As expected, the lower-temperature bottle registers a notable output difference regarding the other two targets, but also it can be confirmed that the sensor is able to discern between background and human (as all but two hand measurements report higher outputs than the maximum background voltage).

Figure 3.36: Measurement of three different targets (mm-wave absorber, human hand, and cold water bottle), where each one has 10 datapoints in the graph. The absorber was at 30 °C, the hand at 33 °C, and the bottle at 20 °C.

3.2.1.2 Image SNR and Contrast

The SNR of the RX throughout the measurement sequence in Fig. 3.34 has been registered and calculated as the traditional imaging-based ratio

$$SNR \text{ (dB)} = 10 \cdot \log_{10} \left(\frac{\mu_{RX}}{\sigma_{RX}} \right) \quad (3.9)$$

where μ_{RX} and σ_{RX} are the mean voltage during the measurement and its standard deviation, respectively. The results may be seen in Fig. 3.37, where the mean is 17.15 dB (with a peak of 19.75 dB and a minimum of 14.83 dB).

Although the passive sensor demonstrates target-separation capabilities in Fig. 3.35 and Fig. 3.36, these are deteriorated by added noise and clutter during imaging scenarios. Noise sources can include the motor's vibrations and the temperature variations inside the imager, while clutter can come from ambient heating or background emissions reflected by the target. In addition, uncontrolled indoor conditions usually offer little contrast in passive imaging,

Figure 3.37: SNR of the passive RX throughout the sequence illustrated in Fig. 3.34. The SNR is computed as the logarithmic scale ratio of the mean signal in each measurement to its standard deviation.

due to warm temperature and strong emissions from nearby objects in the scene. Outdoor scenarios, on the other hand, typically provide better imaging conditions, due to their cooler surroundings, reduced number of nearby clutter, and the possibility of using the cold sky as FoV background for greater contrast.

These considerations force some restrictions on the imager's use in indoor scenarios. Unless otherwise stated, all images are generated in an air-conditioned room cooled to approximately 16 °C to set and control background contrast. Additionally, window blinds in the room are shut to avoid cold reflections from the sky. Finally, nearby objects are removed from the FoV to minimize clutter. The effects of this arrangement can be seen in Fig. 3.38. The raw outdoors image in Fig. 3.38(a) presents great contrast between the subject and the background, whereas in the uncontrolled indoors result in Fig. 3.38(b) the subject is indiscernible. It can be confirmed that the subject's concealment in the uncontrolled scenario in Fig. 3.38(b) is due to surrounding temperature and clutter, which is reduced when the FoV's background is once again the cold sky, as in Fig. 3.38(c). The lower indoor temperature and clutter in Fig. 3.38(d) after following the mentioned restrictions provide adequate contrast to distinguish the subject.

With this conditioning set-up, the passive sensor's raw samples can be directly displayed as an image, but noise prevents accurate visual recognition of objects in the scene. To increase image quality, first a double threshold is applied to limit the represented values within an empirical range¹⁴, filtering some of the remaining noise. Afterwards, a 2-D median filter is run through the image (kernel size is asymmetric, with emphasis on the horizontal dimension, to eliminate column-like noise in the image). Finally, a reduction of the image value interval is performed to enhance contrast. An example of this sequence is shown in Fig. 3.39, for a scenario where a male target wearing a T-shirt hides an IED. The IED is simulated with three cylindrical capacitors totalling 10x15 cm² in size. To offer the best possible images when presenting the processing steps the threat has been cooled to 7 C°; in all other figures the threat is at 16 C° or warmer. Residual noise around the subject's body is unimportant, as a later fusion process with the optical sensor's depth mask erases contributions outside of the body contour [223].

¹⁴The range is manually set during measurements, and depends on the environment (including variables such as clutter, background temperature, type of target, among others), as well as the specific model of the RX.

Figure 3.38: Comparison of raw passive images of a male subject under various atmospheric conditions. In (a) the target is scanned in an outdoors environment, in (b) the target is scanned indoors in uncontrolled conditions, in (c) the target is scanned indoors in uncontrolled conditions in front of an open window facing the sky, and in (d) the target is scanned indoors in controlled conditions.

Figure 3.39: (a) IR camera snapshot of a male target wearing a loose T-shirt and hiding an IED threat (at 7 °C). Colorbar: Temperature in degrees. (b) Raw output image of the passive sensor. (c) Result of thresholding and filtering the raw image with a 2-D asymmetric median kernel. (d) Final passive image, after reducing the image value interval. (e) Graph of a single column in (c), cutting through background, body, and IED. The samples corresponding to the IED can be clearly separated from the body.

It can be noticed that the image quality of the IR camera is superior to the 94 GHz passive RX, due to the camera's better thermal performance and spatial resolution (its thermal sensitivity is $0.05\text{ }^{\circ}\text{C}$ and its wavelength between $7.5\text{ }\mu\text{m}$ and $13\text{ }\mu\text{m}$). On the other hand, many disadvantages arise from working in IR frequencies:

- Transmission through typical dielectric materials used to cover threats, such as textiles, is lower as the working frequency increases [20,224]. This effect can limit the detection capabilities in imaging scenarios with additional clothing layers. For example, if a room-temperature filled water bottle is hidden underneath a thick sweatshirt the contrast between the threat and the body is greater at 94 GHz (as seen in Fig. 3.40), due to greater penetration at this frequency.
- Atmospheric particles (such as humidity, dust, or smoke) produce greater attenuation with increasing frequencies. This effect can limit the detection capabilities in outdoor and/or long-range scenarios. On the other hand, a common atmospheric attenuation window exists at 94 GHz [225].

Figure 3.40: (a) IR camera snapshot of a male target hiding a threat (a filled room-temperature 1 L water bottle) underneath a thick sweatshirt and T-shirt. Colorbar: Temperature in degrees. (b) Filtered and thresholded output image of the passive sensor.

With the use of the IR camera, an approximate state for the ideal performance in the conditioned environment has been found by performing a sequence of measurements in the controlled room as it cools down. The sequence, shown in Fig. 3.41, consists of raw passive images of a male subject with a rectangular target in his chest. The imaging sequence starts with a room-temperature climate, and continues as the air-conditioning is turned on to cool down the surroundings. As can be seen from the results, only after approximately 1 hour of continuous cooling is the ambient temperature low enough to fully make out the subject's body in the raw images. At that instant, the room has been cooled from $25\text{ }^{\circ}\text{C}$ to $19\text{ }^{\circ}\text{C}$. Further cooling shows some improvement: with the room at a $15\text{ }^{\circ}\text{C}$ temperature, the subject's legs can also be detected.

Further testing of the sensor's threat detection capabilities has been performed. A sequence of measurements as a threat grows warmer, to expand on the histogram separations seen in Fig. 3.35, has been performed. A male subject is scanned in a cooled room ($16\text{ }^{\circ}\text{C}$ is the lowest surrounding temperature). The person hides the IED (made of plastic, silicone, and metal components) beneath a T-shirt, having it pressed against his torso. The threat starts the sequence at $4\text{ }^{\circ}\text{C}$ and progressively grows warmer through heat transfer from the

Figure 3.41: Measurement sequence with the passive sensor of a male subject with a rectangular target in his chest, as the room cools down from air conditioning. The sequence shows both 94 GHz and IR images. Colorbar: Temperature in degrees. The image in (a) is taken at the beginning of the sequence, with no air conditioning. In (b) the air conditioning is turned on, and images approximately every 30 minutes are taken in (c) - (f).

subject's body. As expected, the threat becomes smaller when the object warms up (specially seen in the last step). The threat can still be seen in all steps, even as it heats up, because not only is its contrast with the body due to its temperature-related emissions, but also how well it blocks the body's own radiation. Materials more transparent to 94 GHz radiation (such as some plastics) do not prove as detectable as they heat up, since they will allow more of the torso's emissions to pass through. Results may be seen in Fig. 3.42

Figure 3.42: Measurement sequence with the passive sensor as a threat warms up, while hidden beneath a male subject's T-shirt. The threat is the previously-used IED, made up of three individual capacitors. The left column shows IR close-ups of the torso, while the center and right column are the raw and processed images obtained with the passive RX, respectively. Colorbar: Temperature in degrees.

Because the emissions from a concealed target must pass through clothing layers, the effects of these textile materials in threat separation have been analyzed. In a first step, the raw RX outputs when aiming at a cool water bottle are compared in Fig. 3.43, where the bottle is covered with different fabrics in each scenario. The output when the RX aims at a mm-wave absorber is also included for comparison, serving as a background reference.

(a)

(b)

(c)

(d)

(e)

(f)

Figure 3.43: Histograms of the raw RX outputs when the sensor aims at a water bottle in various concealment scenarios, and photographs of some of the set-ups. Outputs from the unconcealed bottle and a mm-wave absorber are also included for comparison. In (a) the concealed bottle is covered with a polyester T-shirt, in (c) with a cotton sweater, and in (e) with denim jeans. The photograph in (b) shows the unconcealed bottle, the one in (d) shows the mm-wave absorber, and the one in (f) shows the covered bottle (with denim).

As expected, the thicker denim creates the worst effect on the bottle's emissions, dispersing the RX's output over values much higher than the unconcealed bottle or even the mm-wave background. In an image, this effect would render the bottle virtually undistinguishable from the background, and less contrasting with the subject's body. On the other hand, the shirt and sweater slightly increase the bottle's output, but still leave enough separation from the absorber to differentiate their histograms. Between these two fabrics, the shirt produces the least noticeable effect; likely due to its thinner material. The corresponding images from each case of the covered or uncovered bottle may be seen in Fig. 3.44.

Figure 3.44: Passive images of a cool water bottle in the concealment scenarios analyzed in Fig. 3.43. In (a) the bottle is uncovered, in (b) the bottle is covered with a polyester T-shirt, in (c) the bottle is covered with a cotton sweater, and in (d) the bottle is covered with denim.

It can be proven that there is margin for the passive images to suffer from additional Gaussian noise while still displaying distinguishable targets. This has been tested by including artificial additive white Gaussian noise (AWGN) in real image data. The original unmodified data can be seen in Fig. 3.45, and the results with added noise are presented in Fig. 3.46, where the noise has zero mean and a standard deviation proportional to the one from the original data in Fig. 3.45(a). As can be seen from the results, even for noise increases as large as 100% of the original data's standard deviation, the subject can still be faintly differentiated. At increases around 150% of the original standard deviation the subject becomes practically undistinguishable from the background noise.

Figure 3.45: Passive image data used for added noise testing. The subject is a male wearing a T-shirt hiding a cool water bottle in the abdomen. (a) Raw samples of the passive sensor's output in digital values, (b) histogram of raw and filtered samples, (c) raw image, and (d) filtered image.

Figure 3.46: Passive image data from Fig. 3.45 with additional AWGN. The left column are histograms of the additional AWGN introduced in each row, the center column is the corresponding raw image, while the right column is the filtered image. The top row is a scenario where the AWGN has a standard deviation equal to 50% of the standard deviation of the data in Fig. 3.45(a), the middle row is a scenario where the AWGN is raised to 100% of the original standard deviation, and the bottom row is a scenario where the AWGN is 150% of the original standard deviation.

3.2.1.3 Cross-Range Resolution

The passive RX's cross-range resolution when placed in the reflector system has been measured. To measure this figure of merit two frozen plastic spheres (2 cm in diameter) have been used as reference targets, held by a human in front of his torso to offer a high-contrast background. The spheres are sequentially brought together in the horizontal and vertical directions, and the minimum separation that allows to distinguish the objects is assumed as the resolution in that direction. The results may be seen in Fig. 3.47, which indicate that the resolution is between 3 cm and 4 cm in both the horizontal and vertical dimensions.

Figure 3.47: (a) Measurement of the cross-range resolution of the passive sensor along the horizontal dimension using the reflector system. 2 cm-diameter frozen plastic spheres are placed 5 cm apart, and sequentially brought together in front of a human. The distances above the images represent the separation between the objects' centers. (b) Same sequence as above, for the vertical dimension.

3.2.1.4 Spot Size

The spot size generated by the passive sensor has also been measured. To do so, a single target from the cross-resolution measurements has been isolated [the left sphere in Fig. 3.47(a)]. A single line scan is taken in the X and Y dimensions, and the sampled digital values are converted first to voltage, and finally to normalized dB. For this sensor the criteria has been a 3 dB cutoff. The measurement may be seen in Fig. 3.48. Results show that the spot is asymmetrical, due to two main reasons. First, since the reflector system's focal point is unique, the multi-sensor configuration dictates that neither the active nor the passive sensor will be perfectly focused (thus producing a deformation of the radiated spot at the standoff distance). Second, the sensors' antennas (horn for the passive, and dipoles for the active), which act as feeders for the imager's reflector system, are also asymmetrical. It must be taken into account that the imager's reflector system comes from a previous commercial design not intended for this hybrid setup, so slight imperfections are expected. A redesign of the imager's reflector system is underway, in order to optimize it for the multi-sensor application.

Figure 3.48: Measurement of the spot size of the passive sensor along the horizontal (X) and vertical (Y) dimensions using the reflector system. A 2 cm-diameter frozen plastic sphere is illuminated, and single line scans in each dimension are measured.

3.2.2 Characterization of the Hybrid Imager's Active Sensor

With the active sensor and support electronics fitted into the reflector system, the initial passive imager is modified to a hybrid system. Just as the passive RX's individual performance has been tested in Section 3.2.1, this section deals with the performance of the 122 GHz radar integrated in the imager. To form an active image, as the reflector system scans the FoV the radar is continuously transmitting an FMCW signal with 150 μ s of chirp time and 15.1 GHz of bandwidth (thus **Configuration 3** from Section 3.1, using only the TRA_120_012 MMIC). Each chirp corresponds to an individual pixel in the 3-D image, where the x and y coordinates are dependent on the position of the scanning mirror and the z coordinate is set as the range measured with the radar. The processing required to form an image is as follows:

- A zero-padded FFT is performed per each chirp period.
- A peak detection algorithm is applied to the spectrum resulting from each FFT. The algorithm selects the range points desired, using a power threshold to control false alarms. More than one range coordinate can be associated to a single pixel, if multiple z layers want to be displayed.
- The resulting 3-D coordinates can be represented in multiple formats, such as a point cloud, an interpolated grid, or a reconstructed surface.

An example of the threat-detection capabilities of the active sensor with the reflector system can be seen in Fig. 3.49, where a male human wearing a T-shirt has been scanned (the same scenario as Fig. 3.39 for the passive sensor). A comparison between the target hiding or not a threat is presented. Photographs of the subject and threat can also be seen. The measured range does not equal the standoff distance (1.7 m) due to the signal propagation through the reflector system. Notice that the three individual columns of the IED are detected in the point cloud representation of Fig. 3.49(b) (blue pixels representing a closer range to the radar, therefore protruding from the subject's torso), and that no such shape is detected when there is no threat hidden. Also of interest is the effect of specular reflection in received power, such as in the vanishing of part of the arms, which is one of the obstacles that can be overcome with the combination of a passive sensor.

Figure 3.49: (a) Photograph of the male subject (the location of the threat underneath the T-shirt is indicated). (b) Scans of the set-up on the left, using the modified active sensor in the reflector system. In the examples on the top the person is hiding an IED composed of three cylindrical capacitors, while on the bottom figures no threat is hidden. Examples on the left are representations in 3-D point cloud format while those on the right depict a surface reconstruction of the point clouds. A photograph of the hidden threat is inset between the examples. Colorbar: Range in meters.

Figure 3.50: Noise power level as a function of the maximum reflected power per pixel for all pixels obtained from Fig. 3.49(b) (with the IED). The red data uses a single range bin near the target to estimate noise power, while the blue data uses the mean power of 50 range bins around the target.

By analyzing the radar's registered noise power in each of the pixels in Fig. 3.49(b) (with the IED) it can be determined that the active sensor is working in a configuration where the noise levels are constant, as illustrated by the results in Fig. 3.50. This means that phase noise is not present around the standoff range, as this would create a dependence on the noise level with respect to the maximum transmitted power (i.e. a non-flat curve in Fig. 3.50) [123]. This is a favorable conclusion because it indicates that the radar's output power could be increased to enhance SNR.

3.2.2.1 Cross-Range Resolution

As with the passive sensor, the cross-range resolution of the active sensor fitted in the reflector system has been analyzed. Two 5 mm-diameter metallic spheres are sequentially brought together in the horizontal and vertical directions. The minimum separation that allows to distinguish the objects is assumed as the resolution in that direction. Steps in the horizontal and vertical sequences can be seen in Fig. 3.51. In both sequences, the targets are still resolved in the central step, which means that the resolutions are symmetric and between 3 cm and 2 cm in both directions. The theoretical diffraction-limited resolution according to (3.3) (at a distance of 2.7 m, to account for the imager's standoff range plus the propagation path in the reflector system) is approximately 1.1 cm, which is much lower than the measured results. This deterioration is most likely due to inefficient illumination of the reflector's aperture: since the commercial imager works with a passive sensor array, the reflector is designed so that each sensor does not fully illuminate the entire aperture to ensure as little spillover as possible when the array moves to scan the FoV.

Figure 3.51: Measurement of the cross-range resolution of the active sensor using the reflector system. The top row is the resolution analysis in the horizontal dimension, and the bottom row in the vertical dimension. The centers of two metallic spheres are placed at a varying distance [(a) and (d) 4 cm, (b) and (e) 3 cm, (c) and (f) 2 cm]. Colorbar: Received power in dBm.

Further testing has been performed to confirm spatial resolution and its degradation throughout the FoV by scanning a matrix of the same metallic spheres seen in Fig. 3.51. These spheres are glued to wooden toothpicks stuck in a polystyrene sheet, and separated 7.5 cm in each direction. The result of this scan may be seen in Fig. 3.52. In the presented area, all pins are accurately separated from each other in both directions.

Figure 3.52: Measurement of the cross-range resolution of the active sensor along the horizontal and vertical dimensions using the reflector system. 5 mm-diameter metallic spheres are placed 7.5 cm apart, in both directions, over a polystyrene sheet. Colorbar: Received power in dBm.

3.2.2.2 Spot Size

As done for the passive sensor in Fig. 3.48, the spot size of the radiated beam can also be measured, by isolating the right-side sphere in Fig. 3.51(a), and the results may be seen in Fig. 3.53. Single horizontal and vertical line scans are shown, centered on the maximum peak of returned power. The distance until that power decreases 6 dB in each dimension is considered the spot size. As with the passive sensor, and for the same reasons, the results show that the spot is not symmetrical.

Figure 3.53: Measurement of the spot size of the active sensor along the horizontal (X) and vertical (Y) dimensions using the reflector system. A 5 mm-diameter metallic sphere is radiated, and single line scans in each dimension are measured.

3.2.3 Sensor Compatibility for Simultaneous Scanning

Although the active and passive sensors operate at different frequencies (122 GHz and 94 GHz respectively), it has been detected during measurements that their simultaneous use can cause interference. The interference is registered as an undesired increase in the passive sensor's raw voltage output. Since only the active sensor transmits radiation, the interference is produced only in the passive sensor.

To quantify this interference an experiment is performed in a laboratory environment, the active sensor is placed outside of the reflector system, aiming directly at the passive sensor just 5 cm away [as seen in Fig. 3.54(a)]. As registered in Fig. 3.54(b), the transmitted chirp causes interferences above the regular output levels, synchronized with the frequency ramp. This is again verified in Fig. 3.54(c), where the contribution from each transmitted frequency in the active sensor's bandwidth to the passive sensor's interference can be seen. Because lower frequencies in the bandwidth are closer to the passive sensor's operating frequency and below the WR-10 waveguide standard's cutoff frequency, these interfere much more. It can be noted that the interference seems to have an asymptotic tendency to null as the frequency increases, meaning that a higher-frequency active sensor could produce no interference.

The effects of the interference caused by the 150 μ s chirp with 15.1 GHz of bandwidth in a passive image can be seen in the comparison in Fig. 3.55. In the imaging case, the interference generated is not only produced by the active sensor's transmission, but also its own heating (which is nearby high-temperature clutter for the passive RX) and its physical blocking of incoming radiation from the scene. The interference causes large spots to appear in the image, masking the target's body. Because both the active and passive sensors are centered in the reflector system, the interference also appears centered in the FoV. This is

(a)

(b)

(c)

Figure 3.54: Interference in the passive sensor's output caused by the active sensor. (a) Measurement set-up, and passive sensor output comparisons in the (b) time and (c) frequency domains.

specially detrimental, because this area usually coincides with the subject's abdomen, where hidden threats are commonly placed. In addition, it can be seen that the interference registers a darker color than the subject's body (confirming its higher-voltage effect on the passive RX's output). The interference can also be seen in the passive sensor's raw output during the image [Fig. 3.55(c)], where spikes saturating the RX's 5 V output can be seen near the middle of the capture.

Figure 3.55: Passive images of a male subject with no hidden threat. (a) Active sensor turned off, (b) active sensor transmitting chirp (interference boxed in red), and (c) comparison of the passive sensor's raw output in both cases.

To solve this interference problem, several alternatives can be studied:

- A naive solution would be to reduce the chirp's bandwidth, and thus transmit less-interfering frequencies, at the cost of deteriorating range resolution. Nonetheless, as the results in Fig. 3.56 show, even when transmitting just 6 GHz (less than half of the original 15.1 GHz bandwidth) the passive sensor's output still presents a slightly higher level (indicating interference) than when no bandwidth is transmitted.
- The sensors can also be separated spatially, instead of in the frequency domain. By shifting the placement of the active sensor in relation with the passive sensor, the location of the interference in the passive image can be moved away from the subject's body. The interference can still be present in the image, but will stop masking threats hidden by the subject. However, this solution is done at the cost of reducing the active sensor's FoV: if the passive sensor is centered in the reflector system and is not seeing the interference in the center of its image, this means that the active sensor is off-center in the array, and thus will not scan the full FoV. This effect can be seen in Fig. 3.57, where the abdomen in the passive image is less masked, but the active sensor is only able to image the subject's arm, shoulder, and ribcage.
- A more complex approach could be the use of a global scene illuminator for the active image, where the active sensor in the reflector system is only tasked with focused reception. The global illuminator could still interfere with the passive RX, but would do so from a farther location, and thus with less effect.
- A simple solution is to multiplex the active sensor's transmission with the passive sensor's reception. This has been the technique chosen for all simultaneous measurements.

Figure 3.56: Dependence of the interference in the passive sensor with the active sensor's transmitted bandwidth.

Figure 3.57: Imaging example when the active sensor is off-center in the reflector system. (a) Passive image with interference shifted slightly away from the subject's abdomen (interference boxed in red), and (b) corresponding active point cloud. Colorbar: Range in meters.

In this time multiplex the passive sensor only captures data when the active sensor is not transmitting, and thus not interfering. To allow these transmission-free instants, the active sensor must be switched from a continuous FMCW operation to a pulsed FMCW transmission. Of course this solution is only feasible in low-speed scans, where the pixel density is high and thus subsampling each sensor's image can be permitted. In the implementation, each FMCW pulse contains a $150\ \mu\text{s}$ chirp with 15.1 GHz, while the off-time between pulses is set to a multiple of the pulse time. For the simultaneous images, the off-time is set equivalent to nine pulses to surpass the passive RX's integration time (although off-times as low as six pulses have also been tested favorably), meaning that the simultaneous images for the active sensor will be subsampled by a factor of ten when compared to the images in Section 3.2.2. Sample images generated with this scheme are presented in Fig. 3.58, while a comparison of the passive sensor output when continuous FMCW transmission is used may be seen in Fig. 3.59¹⁵. As can be noted from the images, the passive sensor now suffers no noticeable interference. On the other hand, the lower pixel density of the raw active image is evident, when compared against the continuous images in Section 3.2.2. Nonetheless, as seen in Fig. 3.58(c), with image processing consisting of interpolation and 2-D filtering, a smoother image with higher quality can be generated. From the graphs in Fig. 3.59 it can be concluded that the pulsed scheme interferes much less than the continuous transmissions, as its histogram presents only a small mean deviation with regard to the data when the active sensor is turned off.

Figure 3.58: Passive and active images of a male subject with no hidden threat when the active sensor transmits pulsed chirps in a 1-to-10 ratio. (a) Raw passive image, (b) raw active image, and (c) interpolated and filtered active image. Colorbar: Returned power in dBm.

¹⁵Note that in the actual implementation, as well as for the images in Fig. 3.58, the sensors are multiplexed so that the passive sensor does not receive as the active sensor transmits. This is in contrast to the data in Fig. 3.59 which shows the passive sensor still being interfered, but this is done simply to demonstrate that the active sensor now has discontinuous chirps in a 1:9 ratio.

Figure 3.59: Passive sensor outputs under different active sensor transmissions. (a) Output in the time domain, and (b) histograms of those outputs.

3.2.4 Image Fusion

In the hybrid imager, the data from the active and passive sensors are output individually in plots like the ones in Fig. 3.39 and Fig. 3.49, which are the traditional results of single-sensor imagers, but it can also be combined to increase the information available in a single representation. For example, a simple combination of the active sensor's and Kinect's outputs could be done by overlaying the two images with a degree of transparency, such as shown in Fig. 3.60. Although such a combination is fast and simple, the information is not intelligently merged: the active sensor's output is not confined to the target's body, and the final result is blurred because of the superposition of both images' pixels.

Figure 3.60: Sample combination of active sensor (a) and Kinect (b) outputs via an overlay (c) of their two images with 50% transparency.

A more effective combination can be performed with image fusion techniques. A simple illustration of the selected combination process is presented in Fig. 3.61. The basis behind the combination is the use of the optical sensor as a body-segmenting tool that provides a depth mask allowing the removal of noise in the active and passive outputs. After fusing with the mask, the active and passive images have their data limited to the inside of the subject's contour. Additionally, the optical sensor provides its depth mask with the texture of its RGB camera. Passive data are not fused with the active sensor to avoid removal of the independent image features of each sensor.

The depth mask extracted from the Kinect's output is critical to the fusion process in

Figure 3.61: Diagram of the image combination process of all three sensors involved in the imager.

order to limit the other sensors' data to the subject's body and remove outlier noise. The creation of this mask is done with the Kinect's SDK without the need for custom-made software, and the representations are done with MATLAB's Computer Vision Toolbox. The Kinect is also able to detect and track the target's skeleton [226]. Thanks to this, a custom face blurring algorithm can be applied if chosen, in order to preserve subject privacy.

While the outputs of the other sensors are being processed, the Kinect photographs the scene with its RGB camera [Fig. 3.62(a)] and performs depth sensing with its IR sensor [Fig. 3.62(b)]. A 3-D point cloud is then extracted from the IR sensor's depth image. Knowing the separation between the color and depth sensors on the Kinect, the RGB information from the color image can then be registered onto the depth point cloud [Fig. 3.62(c)]. This point cloud is then range segmented to isolate the data in the same range window that the active and passive sensors have scanned [Fig. 3.62(d)].

Figure 3.62: (a) Photograph from the Kinect's RGB camera, with face blurring applied and the detected skeleton represented as red segments. (b) Output of the Kinect's IR sensor with a grayscale colormap representing depth. (c) Initial depth point cloud with RGB texture from the camera. (d) Final Kinect 3-D point cloud, later used as depth mask, after segmenting the initial point cloud within the subject's range window.

The outputs from the active and passive sensors are fused with the Kinect mask via 2-D wavelet decomposition of the individual images [227]. The fusion process, for each sensor (active or passive), is as follows:

- The sensor output is saved in a true-color 2-D JPEG format and a discrete wavelet transform (DWT) [228] is applied to the resulting image and the Kinect mask. The transform used is based on Daubechies wavelets [229] with two vanishing moments

(db2) at level 2. The results of applying these transforms to the sensors' outputs are the image decompositions.

- The two sets of decompositions (active or passive, and Kinect) are then combined to form a single batch. The selected merging method is based on choosing the maximum approximations and details, also known as the choose-max technique. By selecting a value from either one image or the other, the end result are the decompositions of a single fused image.
- Finally, the fused image is generated by performing an inverse DWT on the combined decompositions.

From a signal processing point of view [230], wavelet decomposition of an image is simply passing the data through a set of cascading low-pass and high-pass filters. Let $I_1(x, y)$ and $I_2(x, y)$ be the $M \times N$ output images of two sensors. Each image will first be filtered by rows (in parallel, through both the low-pass and the high-pass filter), and then by columns. Thus, for a single image $[I(x, y)]$, a one-level decomposition will produce four results (also called coefficient matrices): $I_{LL}(x, y)$, $I_{LH}(x, y)$, $I_{HL}(x, y)$, and $I_{HH}(x, y)$, depending on the order in which the filters have been cascaded (*i.e.* $I_{LH}(x, y)$ means the image is first low-pass filtered by rows and then high-pass filtered by columns). $I_{LL}(x, y)$ is a smooth subimage that represents a coarse approximation to $I(x, y)$, while the remaining results are the higher-frequency details in the image's vertical, horizontal, and diagonal directions.

If $l(i)$ and $h(j)$ are the low-pass and high-pass coefficients, respectively, of the wavelet used, where $i = 0, 1, 2, \dots, N_l$ (with N_l being the low-pass filter's length) and $j = 0, 1, 2, \dots, N_h$ (with N_h being the high-pass filter's length), then [230],

$$I_L(x, y) = \frac{1}{N_l} \sum_{i=0}^{N_l-1} l(i) \cdot I((2x + i \bmod M), y) \quad (3.10)$$

$$I_H(x, y) = \frac{1}{N_h} \sum_{j=0}^{N_h-1} h(j) \cdot I((2x + j \bmod M), y) \quad (3.11)$$

$$I_{LL}(x, y) = \frac{1}{N_l} \sum_{i=0}^{N_l-1} l(i) \cdot I_L(x, (2y + i \bmod N)) \quad (3.12)$$

$$I_{LH}(x, y) = \frac{1}{N_h} \sum_{j=0}^{N_h-1} h(j) \cdot I_L(x, (2y + j \bmod N)) \quad (3.13)$$

$$I_{HL}(x, y) = \frac{1}{N_l} \sum_{i=0}^{N_l-1} l(i) \cdot I_H(x, (2y + i \bmod N)) \quad (3.14)$$

$$I_{HH}(x, y) = \frac{1}{N_h} \sum_{j=0}^{N_h-1} h(j) \cdot I_H(x, (2y + j \bmod N)) \quad (3.15)$$

A two-level decomposition, such as the one used for this imager, will iterate the filtering on $I_{LL}(x, y)$ to produce another finer-detailed set of decompositions. After this step, the fusion starts by merging the coefficient matrices of the different sensors. The merging technique used for this imager is known as the choose-max method, and is simply the selection of the coefficient with the larger activity level in each image pixel out of the two sensors. The

activity level $[A(p)]$ of a coefficient $[D(p)]$ in image position p from image n can be measured in a variety of ways, but the simplest is to equate it to its absolute value as

$$A_n(p) = |D_n(p)| \quad (3.16)$$

Then, if $Z(x, y)$ is the fused image from combining $I_1(x, y)$ and $I_2(x, y)$, the merging of coefficients is simply

$$D_Z(p) = D_n(p) \quad (3.17)$$

where n is either 1 or 2, depending on which image's decomposition has the highest activity in each position, as expressed by

$$A_n(p) = \max(A_1(p), A_2(p)) \quad (3.18)$$

The fused image $[Z(x, y)]$ is then generated by reversing the cascaded filters over the merged coefficients $[D_Z(p)]$. A diagram summarizing the DWT and fusion process, for the scenario in Fig. 3.49, may be seen in Fig. 3.63.

The image processing and fusion software is run on a conventional PC using MATLAB code. Although processing is not performed in real-time speed, CPU parallelization (via MATLAB's Parallel Computing Toolbox) has been implemented for the processing of the different sensors' data. The complete processing chain is diagrammed in Fig. 3.64. An example of the result of this combination, for the scenario in Fig. 3.39 and Fig. 3.49, is shown in Fig. 3.65. In this case, the Kinect-fused images of both the active [Fig. 3.65(b)] and passive [Fig. 3.65(d)] sensors expose a threat in the abdomen.

3.2.4.1 Additional Active Sensor Image Processing

The initial output for the active sensor, as seen in Fig. 3.49, is a 3-D point cloud where each pixel is colored according to its range or returned power. Before fusing with the Kinect's data, additional processing must be applied to the active sensor's output point cloud. First, the active point clouds are registered with the Kinect point clouds [Fig. 3.62(d)] via the ICP algorithm [170]. This process is performed automatically if the target stands in a pre-determined position. If not, the user must manually input a set of initial reference points in both point clouds for the algorithm to converge to an acceptable solution. The point clouds from the active sensor are then linearly interpolated into a rectangular grid with a sub-cm step size. This grid is then input into a median filter, with a circular kernel with size approximate to the active sensor spot size, to smooth out the representation. Finally, the 3-D filtered grids are projected and saved to a 2-D image format and fused with the Kinect mask through the DWT process described previously.

The various steps discussed are illustrated in Fig. 3.66 (with active images representing both range and received power data), for the imaging scenario in Fig. 3.65. Notice that active images showing returned power information have been included. Although returned power information can sometimes be useful, it is unrecommended as unambiguous confirmation of threat detection in mm-wave frequencies due to its strong incidence angle dependence [123].

Figure 3.63: Breakdown of the processing steps for fusing two images using the DWT-based technique proposed. The unfused input images are from the Kinect and active sensor, for the imaging scenario in Fig. 3.49. Decompositions that include passing through the high-pass filter are reduced to simply the borders of the input image, because these boundaries produce a sharp change in image data (interpreted as a high-frequency component). This is also why the top decomposition, created by passing only through the low-pass filter, is approximately the same image as the input; values inside the image's borders have very smooth changes (interpreted as low-frequency information). The Kinect's depth data has been used to generate image $I_1(x,y)$, but the colored range information in the final fused image $[Z(x,y)]$ comes from the active sensor.

Figure 3.64: Diagram of the complete processing chain and outputs at different stages. Blocks in different colors are processed in separate CPU threads.

Figure 3.65: (a) Imaging scenario of male target hiding a room-temperature IED of three cylindrical capacitors underneath a T-shirt in the abdomen area. In this case both sensors expose the threat. (b) Range-based representation of the active sensor data. (c) Active sensor data fused with the Kinect point cloud. Colorbar: Range in meters. (d) Passive sensor data. (e) Passive sensor data fused with the Kinect point cloud.

Figure 3.66: (a) Range point cloud of the active sensor. Colormap: Range in meters. (b) Registration of the range point cloud with the Kinect's point cloud. (c) Filtered grid with range data. (d) Fused range data and Kinect mask. (e) Returned power point cloud of the active sensor. Colormap: Returned power in dBm. (f) Registration of the returned power point cloud with the Kinect's point cloud. (g) Filtered grid with returned power data. (h) Fused returned power data and Kinect mask.

3.2.4.2 Additional Passive Sensor Image Processing

The passive sensor's output requires significantly less additional processing than the active sensor, because the median filtering and thresholding described in the previous section already reveal any anomalies as empty areas within the body. After this processing, the only additional stage is an automatic colored flood-fill of hole regions inside the torso is performed, in order to highlight possible threats (this is specially useful considering that after the image fusion the red coloring is maintained). As with the active sensor, the flooded image is then registered and fused with the Kinect point cloud. The complete process is shown in Fig. 3.67, for the imaging scenario in Fig. 3.65.

Figure 3.67: (a) Thresholded and filtered output of the passive sensor. (b) Cropped passive image, where the red region corresponds to the result of performing a colored flood-fill of holes inside the torso. (c) Fused passive data and Kinect mask.

3.2.4.3 Imaging Scenarios

The concealed threat scenario of Fig. 3.65 is a basic case for this system: a large, room-temperature object concealed in the warm, predominantly flat and lightly-covered torso area. To ensure that the imager functions properly in complex situations, and to justify the need for the simultaneous use of the active and passive sensors, other scenarios have been tested (summarized in Table 3.7).

Table 3.7: Tested imaging scenarios

Threat	Covered by	Detected by	Presented in
Capacitors	T-shirt	Active & Passive	Fig. 3.65
Pocketknife	Jeans	Passive	Fig. 3.68
Plasticine	Shirt	Active	Fig. 3.69
Bottle	T-shirt	Active & Passive	Fig. 3.70

Fig. 3.68(a) shows the imager's output for a scenario where a male target hides a room-temperature pocketknife in the left pocket of his jeans. This scenario highlights the usefulness of the passive sensor when the active one cannot properly expose the threat. As can be seen from the resulting images, the active sensor is not able to distinguish the threat from the target's clothing [Fig. 3.68(b)], due to two primary reasons: the knife's small volume (compared to the range resolution), and the concealment behind a thick absorbing fabric (denim). On the other hand, due to the difference in temperature and the blocking of

radiation from the subject, the passive sensor can clearly expose an anomaly in the pocket [Fig. 3.68(c)].

Figure 3.68: (a) Imaging scenario of male target hiding a room-temperature pocketknife in the left pocket of his jeans. (b) Range-based representation of the active sensor data. (c) Active sensor data fused with the Kinect point cloud. Colorbar: Range in meters. (d) Passive sensor data. (e) Passive sensor data fused with the Kinect point cloud. This is the only representation in this scenario where the threat is exposed.

Fig. 3.69(a) shows a counterexample of the previous scenario. In this case, the male target is concealing a plastic pouch filled with plasticine and bolts (simulating a plastic explosive with shrapnel) in the stomach area, underneath a long-sleeved dress shirt. In this scenario the threat has been in contact with the human skin during an extended period of time, and thus the plasticine has almost reached body temperature. Because the temperature difference is smaller than the passive sensor's resolution it is not able to expose the threat [Fig. 3.69(c)]. On the other hand, the active sensor can accurately expose a protrusion from the target's stomach [Fig. 3.69(b)]. In this situation, the combined use of both sensors enables threat exposure when the passive sensor alone is not sufficient.

Figure 3.69: (a) Imaging scenario of male target hiding a simulated plastic explosive with shrapnel, which has warmed up to body-temperature, in the right side of his abdomen. (b) Range-based representation of the active sensor data. (c) Active sensor data fused with the Kinect point cloud. Colorbar: Range in meters. This is the only representation in this scenario where the threat is exposed. (d) Passive sensor data. (e) Passive sensor data fused with the Kinect point cloud.

Lastly, Fig. 3.70(a) shows a scenario where again both sensors are able to expose the threat. In this case, the male human target is concealing a room-temperature water-filled 1 L plastic sport bottle in the stomach area, underneath a T-shirt. The active sensor can accurately expose a bump along the target's stomach [Fig. 3.70(b)]. Additionally, since the water has not yet reached a temperature equilibrium with the body, the passive sensor can also expose the object [Fig. 3.70(c)].

Figure 3.70: (a) Imaging scenario of male target hiding a plastic 1 L sport bottle full of room-temperature water underneath a T-shirt in his stomach area. In this case both, sensors expose the threat. (b) Range-based representation of the active sensor data. (c) Active sensor data fused with the Kinect point cloud. Colorbar: Range in meters. (d) Passive sensor data. (e) Passive sensor data fused with the Kinect point cloud.

3.2.5 Array Scanning

The imaging examples presented in the previous sections are generated using a single sensor of each class (active, passive, and optical). Nonetheless, the imager's pedestal has enough space to fit several sensors in an array configuration. When multiple sensors of a single class are placed side-by-side in the reflector system, various areas within the FoV can be scanned simultaneously, as illustrated in Fig. 3.71(a). This allows reducing the length of the pedestal's lateral movement, since to cover the complete FoV the pedestal now only has to perform a movement equal to the distance between sensors in the array. As a consequence of reducing the platform's displacement, the single-sensor scan speed (v_{single}) is increased by

$$v_{multiple} = N_{sensors} \cdot v_{single} \quad (3.19)$$

where $v_{multiple}$ is the scan speed with multiple sensors in an array, and $N_{sensors}$ is the number of sensors in the array.

Figure 3.71: (a) Diagram of array scanning with two sensors, placed on a moving platform, simultaneously transmitting at different parts of the FoV. (b) Photograph of the four passive RXs mounted in an array configuration in the reflector system.

With this benefit in mind, the imager has been fitted with several sensors in its pedestal, to generate images with an array configuration. For the first scans, four passive RXs are placed in the imager, as shown in Fig. 3.71(b). In this set-up no active sensors are used, so only a passive image is generated. A sample scan in this configuration, when the target is a human male hiding no threats, is presented in Fig. 3.72. In Fig. 3.72(a) the raw scans of the four RXs are concatenated horizontally, as can be seen by the slightly different colormaps. It can be noted that the spacing between the sensors in the array is such that there is some overlap between their partial FoVs, due to the fact that some of the subject's body parts are duplicated. To eliminate this effect, in Fig. 3.72(b), overlapping columns in the raw image are erased. In addition, the median filter described in Section 3.2.1 is run through the image to eliminate noise. Although the RXs used are of the same model, their mean raw outputs are not exactly equal, causing the notable differences in Fig. 3.72(b) background's color. This can be further seen in Fig. 3.73, where the raw outputs of the four RXs (RX1 - RX4) during the scan in Fig. 3.72 are compared. To solve this issue, the mean of every RX's data in Fig. 3.72(b) is normalized. Afterwards, the complete image data are thresholded (as done in Section 3.2.1), with the result seen in Fig. 3.72(c). In future multi-sensor tests, a pre-scan calibration procedure to compensate this effect should be established.

Figure 3.72: Multi-sensor scanning of a male subject with four passive RXs in an array configuration. (a) Joined raw outputs of the 4 sensors, (b) filtered outputs, and (c) normalized and thresholded outputs.

The active sensors can also be placed in an array configuration. At the time of these tests only two 122 GHz radar boards were available, and thus only two active sensors can be used simultaneously. Their placement in the pedestal may be seen in Fig. 3.74, where updated reduced-size versions of the radar have been used (Annex D). The imaging result when scanning a human male hiding no threat with this array set-up is presented in Fig. 3.75. Because only two sensor are used in this scenario, the pedestal still needs to move a considerable length to cover the entire FoV. Consequently, the partial FoVs of each sensor already cover almost the entire human subject (except one arm each). Just as with the passive RXs, the active sensors present slight differences in noise levels, which means that for a fixed detection threshold the background noise in their images varies. This is why the image from the second

Figure 3.73: Histograms of the raw outputs of the four passive RXs during the imaging scenario in Fig. 3.72.

active sensor [in Fig. 3.75(b) and (e)] presents less background noise than the one from the first sensor, even though both have the same -70 dBm cutoff. If the individual sensor images are combined as they are [such as done in Fig. 3.75(c) and (f)], the result maintains the noticeable difference in background noise. To solve this problem, different detection thresholds are used for the processing of each sensor's images before they are combined. The results of this sequence can be seen in Fig. 3.76.

Figure 3.74: Photographs of two 122 GHz radars placed in an array configuration for simultaneous scanning in the imager.

In future work, simultaneous scans with multiple sensors of both types should be performed. Due to restrictions from the pedestal's size, initial scans would be limited to only two sensors of each type in a horizontal array. As a solution, multi-height arrays (where sensors are stacked on top of each other) could be tested.

Figure 3.75: Multi-sensor scanning of a male subject with two 122 GHz radars in an array configuration. a) Point cloud from the first sensor, (b) point cloud from the second sensor, (c) and combined point clouds. Colorbar: Range in meters. (d) Point cloud from the first sensor, (e) point cloud from the second sensor, and (f) combined point cloud. Colorbar: Returned power in dBm.

Figure 3.76: Combined multi-sensor scans of a male subject with two 122 GHz radars in an array configuration, after processing each sensor's images with different detection thresholds. (a) Colorbar: Range in meters. (b) Colorbar: Returned power in dBm.

3.3 Future Work

This section details some of the research lines and future work to improve the developments of this chapter. Improvements for the 122 GHz and 94 GHz hybrid imager presented in the previous sections are described.

Because the hybrid imager is still a work-in-progress prototype, several improvements can be made to enhance performance in multiple areas. Due to its imaging application, the topmost priority for future work will have to be improving resolution. Although the passive images could also be upgraded, the active images are the ones with the more urgent need of refinement in the cross-range and range dimensions, specially when compared against other state-of-the-art active imagers (such as the 300 GHz system presented in Section 2.1).

In the range dimension, increased resolution requires a larger transmitted bandwidth, which can only be achieved by substituting the current active sensor for another device. In cross-range, the resolution could be increased by operating at higher frequencies and/or optimizing the imager's reflector system for the new hybrid set-up. To operate at a higher frequency, some solutions (in the form of new MMICs at 245 GHz) are being commercially developed by Silicon Radar GmbH [231]. Alternatively, a higher operating frequency could be achieved by using the TRX chain of the 300 GHz radar presented in Section 2.1 as the hybrid imager's active sensor. To this end, two set-ups are possible. The 94 GHz RX could be mounted in the 300 GHz radar's reflector system [as shown in Fig. 3.77(a)], or viceversa. Both of these set-ups have range-associated difficulties: the 300 GHz reflector system is designed to achieve a 8 m standoff distance, for which the passive RX has not been tested, while the 94 GHz reflector system is designed for a much closer range filtered out in the TRX chain's IF path. As an added obstacle, interferences from the 300 GHz chain in the passive RX have also been encountered [presented in Fig. 3.77(b)]. It is worth noting that the interference at 300 GHz are substantially lower than at 122 GHz. According to Fig. 3.77(b), the chirp at 300 GHz creates a 1.105 V mean output in the passive RX, just 1.75% higher than the 1.086 V output with no chirp. Meanwhile, according to Fig. 3.54(c), a 122 GHz chirp produces a 2.358 V mean output in the passive RX, 117.12% higher than the no-chirp output.

Figure 3.77: (a) Set-up for hybrid imaging with the 94 GHz RX mounted in the 300 GHz reflector system. (b) Interferences in the 94 GHz RX caused by the 300 GHz radar's TRX.

A simpler alternative to enhance cross-range resolution could be redesigning the imager's reflector system to take into account the new hybrid multi-sensor array implemented. In this regard, a collaboration with the *Grupo de Antenas, Radar, y Comunicaciones Ópticas* of the *Universidad de Vigo* has been initiated, to first analyze the current reflector system, and from then on detect possible optimizations when simultaneously feeding the system with

active and passive sensors. Initial results show that with the current reflector system spot sizes as low as $1.2 \times 1.1 \text{ cm}^2$ are achievable at 122 GHz ($<50\%$ the current size) by placing the active sensors at a different distance from the reflector and shifting the standoff distance.

Aside from image quality, other critical improvements have also been planned:

- To remove the clutter, multi-path reflections, and FoV-blocking created by the active sensor's PCB, a new reduced-size board currently in development can be used (Annex D). Among other improvements, it reduces the sensor's area from $150 \times 150 \text{ mm}^2$ to just $50 \times 50 \text{ mm}^2$, producing a noticeable improvement in the active images generated, as seen in Fig. 3.78 and harnessing the full $2 \times 1 \text{ m}^2$ FoV without vanishing near the legs.

Figure 3.78: (a) Nude female mannequin used as target, (b) active image of the mannequin with the previous PCB, and (c) active image for that same scenario with the new PCB. Colorbar: Returned power in dBm.

- To increase weapon-detection accuracy, automatic threat recognition (ATR) algorithms and improved fusion methods are being considered. Although the current image fusion technique generates qualitatively good images for human-based recognition, more complex fusion approaches can highlight additional image features specifically optimized for machine learning algorithms. To this end, ATR approaches have already been tested in similar images generated from active and passive sensors in a passenger-screening application, as is detailed in Chapter 4. In addition, an extensive measurement campaign with various threats and subjects is being performed, to better characterize each sensor's threat-detection performance against different materials and clothing.
- To boost processing speed in an effort to develop real-time capabilities, the current image processing code is being transferred to a GPU platform using the experience from Section 2.2. If implemented along with a multi-sensor array in the reflector system (to enable the fastest available mechanical scanning in the imager, at 1 fps), real-time processing software would allow an increase in passenger throughput. To this end, the image fusion stage has already been implemented in CUDA code, with the use of custom-made software based on a publicly-available GPU implementation of the DWT [232]. For this fusion stage and the images in Fig. 3.65, the parallel software allows reducing execution time to just 0.34% of the serial version, with very little error¹⁶ in the borders stemming from approximations and truncations performed in

¹⁶Just 1.61% mean normalized pixel difference between GPU-processed and CPU-processed images.

the GPU code.

- To enable multi-frequency passive scanning (as done in other state-of-the-art hybrid imagers [92]) and generate passive images with higher resolution, the IR camera used for testing can be attached to the imager. Together with the 94 GHz sensor, this addition would provide passive images at two different frequencies, which allows a more thorough characterization of a target through analysis of its radiation in two different bands. In addition, the higher quality of the IR camera's images could ease the computation of automatic detection and recognition algorithms. Because the camera is also a COTS device, the low-cost approach to this imager's design is not affected.

Chapter 4

Automatic Threat Detection (ATD)

This chapter deals with the development of a machine learning model for automatic concealed threat detection based on neural networks. The model is applied to both active and passive images generated with systems in the mm-wave and IR bands. The active mm-wave images come from a public dataset created for airport security, while the collection of passive IR images has been generated in-house. The results show the feasibility of these models to automatically detect hidden threats in multi-sensor checkpoints. The future goal of this development would be the transfer of these models to the 300 GHz radar and hybrid imager described in the previous chapters.

The following sections are included in this chapter:

- Section 4.1 presents the threat detection model for active images generated with a mm-wave airport scanner. The dataset is described and presented with some sample images, after which the neural network architecture is explained. Lastly, the results of applying the developed model to the image dataset are presented.
- The knowledge gained from the processing of active images is applied to a passive IR dataset in Section 4.2. This section follows the same structure as the previous one, but also introduces variations to the network architecture to account for the differences in input images.
- The chapter ends in Section 4.3, which summarizes future improvements for the developed models.

4.1 ATD in Active Millimeter-wave Images

A critical aspect in improving passenger throughput through security checkpoints is the speed at which the resulting images are inspected for possible threats. In this regard, even with state-of-the-art high-speed scanning, bottlenecks can be created if the threat detection procedure is slow [99]. Traditional passenger and luggage screening used human operators to determine whether or not images contain any threats. Current security systems use a mix of human verification and machine learning approaches to automatically detect anomalies. In this scheme an automatic threat detection (ATD) algorithm is run on the scanner's images, resulting in potential threat areas highlighted on a display for the trained human operator to

verify, as shown in Fig. 4.1. Probable areas can then be further inspected manually via pat-downs or hand-held detectors. With additional steps and complexity included to the ATD algorithm, and if the image contains enough information, threat recognition (i.e. classifying the specific anomaly detected within types of threats) can also be accomplished [233]. As ATD algorithms increase their accuracy, the need for human verification and intervention is reduced, ultimately trending towards a fully-automated threat detection system which can screen passengers as fast as the imager can scan them.

Figure 4.1: Sample display from a mm-wave security scanner from L-3 Technologies running ATD algorithms.

In addition, automated detection offers another advantage in the form of lessening the privacy concerns associated to the inspection of a passenger’s clothes-penetrating images with human operators. The careless exposition or storage of these images have generated public backlash causing major setbacks to the spread of higher-resolution scanners worldwide [234–236]. ATD approaches can solve these concerns by processing the revealing images without displaying them, and only outputting a threat probability, list, or simplified silhouette.

Within the framework of this thesis, ATD algorithms would be useful to both the 300 GHz radar and the hybrid imager. For this purpose, a prototype ATD algorithm based on convolutional neural networks (CNNs) has been created. The algorithm receives a passenger image from the scanner and outputs threat probabilities for different body areas present in the image. Although the algorithm must be trained beforehand (a procedure which can require many hours of computing time), new images can then be tested in real-time due to the use of GPU hardware. To appropriately train the ATD algorithm, a dataset of images is required. The larger the dataset, the more varied the scenarios tested, resulting in higher accuracies with new images. With this in mind, photographic datasets such as ImageNet [237], used to train state-of-the-art networks for tasks such as image classification and object localization, contain millions of images with thousands of different categories. Because no in-house dataset of such volume was available for the 300 GHz radar and hybrid imager, the images used to prototype the developed network have been obtained from a public source using scans generated in the 24 GHz - 30 GHz band.

4.1.1 Image Dataset

The ATD system developed has been tested with images from an airport security screening dataset created by the DHS and the TSA, of the United States of America. The dataset was made publicly available on 2017 through the Kaggle platform, an online community owned by Google where data science and machine learning competitions are regularly held. The DHS provided the dataset under the *Passenger Screening Algorithm Challenge* [238], where online contestants were given 6 months to develop an algorithm able to detect concealed threats

in the provided images. The motivation behind this competition, as expressed by the TSA, was “challenging the broader data science community to help improve the accuracy of the TSA’s threat prediction algorithms”. According to the TSA, this improvement is needed to reduce the number of manual screenings obligatory when current threat detection algorithms produce false alarms, because these create significant bottlenecks at security checkpoints. For competitors, the goal in this challenge was to develop an algorithm that detects threats in the dataset with one of the best eight accuracies, in order to win up to \$0.5 million in prize money. With 150 participating teams, although our developed algorithm did not place within the final top eight entries, it was consistently within the top ten solutions of the competition until the final stages.

The dataset provided contains 3-D full-body scans of multiple passengers generated with a mm-wave scanner the TSA calls High Definition-Advanced Imaging Technology (HD-AIT). Although not explicitly stated, the HD-AIT system is believed to be L3’s ProVision system [101], which operates between 24 GHz and 30 GHz. The provided images include scenarios with multiple body types and ages, both genders, various kinds of clothing, and several threat classes and placements, for a total of more than 3 TB of data. In Fig. 4.2 sample images from the dataset are presented.

Figure 4.2: Sample images from the DHS dataset. (a) Front view of a man with no threat. (b) Side view of a woman with no threat. (c) Side view of a man with a concealed threat near his right knee. (d) Side view of a woman with a concealed threat near her left calf. (e) Side view of a man with a concealed threat near his right ribcage. (f) Back view of man with a concealed threat near his shoulder blades. Threats boxed in blue, only for this display.

As detailed in the dataset’s information, from the prominent glint in the samples presented it is apparent that the provided images contain the scanner’s amplitude data, and not its range information. This is to be expected from an imager such as the ProVision system, which operates with a reduced bandwidth (meaning low range resolution) but with multiple TRXs in arrays (allowing great cross-range resolution). Although the images are of returned

signal amplitude in volts, the absolute scale is not provided. As can also be noticed from Fig. 4.2, the dataset contains images from multiple viewpoints around static persons holding their arms above their heads, as usual in portal-based scanners like the ProVision system. In fact, for each person scanned, the dataset contains 16 different images (each from a different angle) forming a full 360° view around the subject. A sample sequence of a full rotation around a subject is shown in Fig. 4.3.

Figure 4.3: Sample 360° scan of a male subject in the DHS dataset.

With the images provided in the dataset, the challenge consisted in detecting any hidden threats carried by each subject in up to 17 different body zones covering the full body surface¹. The body areas considered, and their numbering as used throughout this thesis and the challenge, are depicted in Fig. 4.4. In addition, each subject can be carrying more than one threat, each in a different body zone. Threats in the dataset come in different shapes and sizes, and are not particular to a specific body area only. Notice that because no single body area is visible in all of the frames of a full 360° rotation (like the one in Fig. 4.3), a threat can be occluded in a specific frame and visible in another.

¹Except the subject's head, where threats would be much harder to hide.

Figure 4.4: Diagram of possible threat locations for the images in the DHS dataset. Reproduced from [238].

For each subject scanned, files in the dataset are provided in four formats:

- **.aps**: Projected image angle sequence file. This is the sequence of 2-D equally-spaced snapshots around the subject shown in Fig. 4.3. Each file has a size of 10.3 MB.
- **.a3d**: Combined image 3-D file. Composite 3-D image formed from equally-spaced frames around the subject. Each file has a size of 330 MB.
- **.a3daps**: Combined image angle sequence file. Sequence of 3-D frames for a fixed angular increment around the subject. Each file has a size of 41.2 MB.
- **.ahi**: Calibrated object raw data file. Raw captured data from the scanner, from which the other three data types are generated through various geometric projections. Each file has a size of 2.26 GB.

For the solution implemented only 2-D images have been used, in order to develop an algorithm as least computationally-intensive as possible to run in real-time. The dataset is composed of 1247 labeled scenarios and 100 unlabelled cases, each with 16 images corresponding to the different viewpoints. The labeled data inform whether or not the subject is carrying a threat in each body zone, in order to provide ground truth for algorithm training. The unlabelled data contain no information regarding threat existence or placement, and is thus used to test the developed algorithm with unknown images. The result of these tests were to be submitted to the Kaggle platform in a .csv format listing each image's identification numbers and the normalized probability of threat existence in each body zone. Kaggle then compounds the individual probabilities for all images into a single score (S) according to the binomial cross entropy loss

$$S = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] \quad (4.1)$$

which measures the dissimilarity between image labels and their predictions, where N is the number of scan scenarios times the number of body zones, y_i equals 1 if a threat is present in the given body zone and 0 if not, and $\hat{y}_i$ is the predicted probability of threat existence in the body zone. The lower the loss, the more similar the labels and their predictions, indicating better learning and performance. The final score achieved by our solution was $S = 0.16795$, finishing in the top 40 submissions.

4.1.2 Threat Detection with a Convolutional Neural Network (CNN)

Compared to photographic data, little academic literature is available regarding machine learning for submm- and mm-wave security imagers. This is most likely due to the comparatively low quality of the resulting images, which handicaps an algorithm's learning process. Within the existing literature the majority of examples are related to passive scanners, because these have been widespread for a longer time. The examples can then be subdivided into two major categories, depending on the type of machine-learning approach used: an algorithm relying on traditional hand-engineered stages to extract specific features from the image, or a deep learning technique using a neural network to organically learn from the images. Once again, because hand-engineered approaches have been available before deep learning methods, most publications make use of the former. To extract image features and detect objects, common hand-engineered approaches for passive scanners use principal component analysis [239], local binary patterns [240–242], Fourier descriptors and other contour-based approaches [243, 244], as well as various segmentation strategies [245–248]. Although hand-engineered methods are more widespread, deep learning approaches for passive scanners have also been presented [249]. In a similar manner, although the number of publications related to passive scanners is larger, examples of machine-learning for active imaging are also available [250–252]. Outside of the mm-wave frequency range, several machine-learning publications applied to object detection in security screening for the visible [253], IR [254–256], and X-Ray [257] spectrums have also been presented.

When comparing the two different machine-learning approaches presented, it must be emphasized that hand-engineered approaches rely on expert knowledge regarding the dataset and application in order to successfully design the feature extraction method. In contrast, a pre-developed neural network with demonstrated performance for detection in a specific dataset can be easily modified to perform a similar task in a different dataset. For a developer with a relatively low computer vision background, this allows adapting already-designed state-of-the-art networks in virtually no time. The tradeoff is the large requirement in computing power needed to quickly train a neural network, but the availability of the GPU hardware used for the imaging radar in Section 2.2 allows bypassing this issue. The optimization of GPUs for matrix operations (necessary when processing image inputs) provides large reductions in training duration. For these reasons, a neural network has been the chosen framework for the developed detection algorithm.

The fundamental aspects of neural networks can be traced back to the 1950s with experiments analyzing the functioning of animal visual cortexes [258]. In these measurements, the authors detected varying activity in an animal's neurons depending on the visual stimuli presented. Hence, the neurons were determined to be acting as filters, producing different patterns for different visual inputs.

In deep learning, neurons can be most simply characterized by what is called a perceptron [259]. This algorithm, shown in Fig. 4.5, takes a group of inputs ($x_1 \dots x_m$) and multiplies them by a set of weights ($w_1 \dots w_m$). Afterwards, bias (b) is added to the result before passing through an activation function [$f(\sum)$]. The function produces a single output (φ), and commonly incorporates a non-linear aspect. Without this non-linearity, the perceptron would be reduced to a series of multiplications and sums, producing a simple linear transformation of the inputs, which is not appropriate for complex learning. Biologically, these non-linear functions simulate the activation (or firing) of neurons in the cortex. The operations performed by the perceptron can then be summarized as

$$\varphi = f\left(\sum_{i=1}^m x_i w_i\right) \quad (4.2)$$

Figure 4.5: Flow diagram of a perceptron.

The single-layer perceptron can be cascaded to form a multi-layer neural network, as done in the pioneering work of LeCun [260, 261]. This work has served as a foundation for the development of CNNs; a neural network where the operation performed by each neuron is a convolution. Compared to the initial multi-layer perceptrons, CNNs present several advantages, such as reduced dimensionality and forced connectivity between neurons, leading to better generalization during learning. Fig. 4.6 shows a simplified diagram of a generic CNN as illustration. In this example, the network receives an RGB image (3 channels) with a 128x128 pixel size. A first convolutional layer outputs a 64x64x6 volume, after which a max pooling layer reduces the dimensions to 16x16x12. Finally, the last layer outputs a 1x256 vector which provides some kind of summarized information about the input image.

Figure 4.6: Generic diagram of a simple CNN with few layers.

In CNNs (and many other networks) the learning is done thanks to the use of backpropagation [262], which is the backwards propagation of error through the network. When a labeled input passes through the network, the output's error with respect to the expected result can be computed. Before the next set of inputs passes through, this error is then propagated backwards to the layers, so that the weights can be appropriately updated to reduce error.

Although the success of LeCun and others with CNNs for recognition tasks was proven early, it was not until 2012, with the creation of the AlexNet [263] architecture, that CNNs began attracting widespread attention. AlexNet demonstrated the feasibility of training deep CNNs for very precise recognition with large databases (in part thanks to the concurrent rise and availability of GPU hardware for computing). The network was able to classify more than 1 million ImageNet images into 1000 different classes with a 15.3% error rate (10.8 percentage points below the runner-up and 9.6 percentage points lower than the previous year's best algorithm), vastly outperforming traditional hand-engineered machine learning algorithms. AlexNet is composed of 15 layers, including 5 convolutional layers, 2 normalization layers, 3 pooling layers, 2 dropout layers, and 3 fully connected layers. The simplified architecture can be seen in Fig. 4.7.

Figure 4.7: Commented diagram of AlexNet’s architecture, based on an illustration in [263]. Notice that the network’s flow has two stacked layers dividing processing of the input image in two halves, due to the original authors’ use of two GPU’s during development.

Each of the layers in AlexNet has the following function:

- **Convolutional layer.** In this layer each neuron applies a convolution operation using a set of weights (called filter) to a region of the input whose size is determined by the designer. In the case of an image input, the structure of the neurons in the layer is such that the convolution is ultimately applied the full image. With image inputs, the receptive field (area of the input where each neuron applies its operation) is usually small, with common values of 3x3 or 5x5. To decide which neuron in the layer is activated after processing the input, an activation function is used. In AlexNet this function is the rectified linear unit (ReLu) function [264] which computes

$$f(x) = \max(0, x) \quad (4.3)$$

outputting the node’s result only if its positive.

- **Normalization layer.** In this layer its inputs (usually the outputs of convolutional layers) are normalized among the batch of images that was just processed through the previous layer. This operation is done to increase accuracy by avoiding overfitting the model to each batch and providing better generalization.
- **Pooling layer.** With this layer the size of the network is reduced by pooling a group of its inputs into a single output. This is usually applied after a convolutional layer (with a normalization layer in between) to combine its outputs. AlexNet uses a max pooling strategy, which selects the maximum from a number of inputs defined by the designer.
- **Dropout layer.** This layer performs a deliberate disconnection of node outputs with a fixed probability determined by the designer. This avoids always training all data with all the network’s nodes, forcing the remaining ones to learn more robust image features and thus increasing accuracy. In addition, the passes through the network are sped up due to the reduction of node interactions.
- **Fully connected.** This layer connects each input to all outputs, thus functioning as a convolutional layer with a filter size of one where each neuron takes the full input set as receptive field. The connection can be done with linear or non-linear functions. In

AlexNet the first two fully connected layers also use the ReLu function, while the last layer uses a softmax function [265]. This is because the last fully connected layer is also the last network layer, which must finally classify the input image. The softmax function converts an input vector into a series of normalized probabilities for all possible image scenarios accounted for by the network.

For simplicity, network design, training, and testing have been implemented using MATLAB's Deep Learning Toolbox. The network used has been a modified AlexNet CNN, pre-trained on the ImageNet dataset. Using a pre-trained network is a common technique called transfer learning [266], which allows faster training of new data, as the network has already learned several image features and converged to a stable model. The network can then simply be fine-tuned to the new data with a faster and easier training process using a known starting point for the model. To set the starting point, the weights and biases of each layer are simply set to the values computed after the pre-training. Transfer learning is usually a prerequisite when training networks with small datasets, as is the case with the DHS collection. Although ImageNet images are very different to the DHS dataset, some of the general image features learned can still be applied.

The pre-trained network is then modified to adjust to the new input dataset. For this, the first step is to modify the size of the last fully connected layer from 1000 possible ImageNet classes to just 2 (i.e. whether a threat is present or not). Because in a CNN the initial layers tend to learn only generic image features, such as edges or shapes, while the final layers learn task-specific information, only these latter layers must be changed in the pre-trained model. In our implementation, the three fully connected layers have been modified, as well as the final convolutional layer. The modification comes in the form of changing the weights and biases applied for the convolution operation in each layer. The pre-trained values are changed to randomized magnitudes² following a normal distribution, allowing the network to forget the ImageNet-specific information. In addition, the learning rate factor (the speed at which these weights and biases are updated during training) is increased for the modified layers in order to allow them to catch up to the converged state of the unmodified layers. A summary of all layers involved in the implemented CNN model may be seen in Table 4.1.

The network's parameters during training were set to the following values:

- **Momentum = 0.9.** This parameter is used to specify how much the CNN's previous training iteration contributes to the current iteration's updates. The parameter takes values from 0 to 1, with higher values indicating more contribution. Its physical meaning is similar to that of friction in a dynamic system, which prevents large and sudden movements, while simultaneously encouraging constant progress in a single stable direction.
- **Initial learning rate = 0.001.** This factor determines the rate at which the CNN learns, by adjusting the magnitude of the weight update during backpropagation. Such a low value is due to the use of a pre-trained network. Because the deeper layers are modified for increased learning rate, this value only affects the initial layers³, which have already reached an optimal state during the ImageNet pre-training and thus little to no updates are desired.
- **Batch size = 256.** This is the number of images included in each forward pass of the network, after which weight updates are performed. To speed up training this value was set to the maximum available considering the GPU's memory. On the other hand,

²Setting to zero is usually not a favorable starting point because it can nullify the weights of connected layers during updates in backpropagation.

³All layers except the modified ones, which are the three fully-connected layers and the last convolutional layer.

batch sizes too large will cause abrupt learning curves, as large loads of information are passed through the network in each update.

Table 4.1: Summary of the network architecture used, based on a modified AlexNet CNN

Layer [†]	Type	Activation Size	Weight Size	Bias Size
1	Image Input	227x227x3	N/A*	N/A*
2	Convolutional	55x55x96	11x11x3x96	1x1x96
3	ReLu	55x55x96	N/A*	N/A*
4	Normalization	55x55x96	N/A*	N/A*
5	Max Pooling	27x27x96	N/A*	N/A*
6	Convolutional	27x27x256	5x5x48x256	1x1x256
7	ReLu	27x27x256	N/A*	N/A*
8	Normalization	27x27x256	N/A*	N/A*
9	Max Pooling	13x13x256	N/A*	N/A*
10	Convolutional	13x13x384	3x3x256x384	1x1x384
11	ReLu	13x13x384	N/A*	N/A*
12	Convolutional	13x13x384	3x3x192x384	1x1x384
13	ReLu	13x13x384	N/A*	N/A*
14	Convolutional	13x13x256	3x3x192x256	1x1x256
15	ReLu	13x13x256	N/A*	N/A*
16	Max Pooling	6x6x256	N/A*	N/A*
17	Fully Connected	1x1x4096	4096x9216	4096x1
18	ReLu	1x1x4096	N/A*	N/A*
19	Dropout	1x1x4096	N/A*	N/A*
20	Fully Connected	1x1x4096	4096x4096	4096x1
21	ReLu	1x1x4096	N/A*	N/A*
22	Dropout	1x1x4096	N/A*	N/A*
23	Fully Connected	1x1x2	2x4096	2x1
24	Softmax	1x1x2	N/A*	N/A*
25	Classification Output	1x1x1	N/A*	N/A*

[†]Layers are presented in sequential order.

*Not applicable. The layer does not include weights or bias.

- **Number of epochs = 1000.** An epoch is completed when a forward and backward pass through the network has been performed for the full dataset. In our implementation, this value presents an upper limit only, since the network is conditioned to stop the learning stage once it reaches an accuracy goal. This technique, called early stopping [267], avoids overfitting in cases where the network has already reached a stable state, and subsequent epochs will only restrict the model to the specific inputs. For the developed model, the early stop condition is reached when the network achieves an average accuracy of 99% over the last 100 iterations.
- **Learning rate drop factor = 0.9.** This is a multiplicative factor applied to the weight and bias learning rates to lower its value as the training progresses. The factor takes values between 0 and 1, where smaller values decelerate the network's updates more during training. The learning rate is multiplied by this factor once a certain number of training epochs have passed, defined by the learning rate drop period. This technique is done to avoid sudden variations in the model in the later stages of training, where the network is supposed to smoothly be converging into a stable state.
- **Learning rate drop period = 10.** This is the number of training epochs until the learning rate is reduced via multiplication with the learning rate drop factor. This parameter must be fitted to the specific training duration of each network; setting it

with a value higher than the number of epochs which the CNN needs to reach a stable state will be useless.

- **L2 regularization = 0.01.** L2 regularization uses the sum of squares of all weights to measure a model's complexity, under the premise that the higher the weight values the more complex the network. This is done because a more complex network can usually indicate overfitting to the training data, instead of a good generalization of less specific image features. Through the complexity quantification of L2 regularization, more complex networks can be penalized during training to avoid overfitting. The parameter here is a multiplicative factor applied to the sum of weight squares in the loss computation stage of the CNN's training to specify how much the network takes its complexity into account during training. Higher parameter values will take weight values more into account when performing layer updates.

To facilitate the learning process, the developed model is actually composed of 17 different networks. In this scheme, each CNN is tasked with performing threat decisions on a single body area. Although all networks are fed the same full images, this organization allows each network to focus and learn features from a reduced section of the image. The tradeoff to this structure is the increase in the duration of the full model's learning, which must now wait for 17 networks to reach their accuracy goal.

In relation to this structure, because the frames with different viewpoints in each scan scenario are not labeled individually, it is not known in which of them the target is visible. To bypass this issue a single scenario is studied to subjectively decide which body areas are visible in each frame. From this analysis a matrix M of visible areas for each frame is extracted as

$$M = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \end{bmatrix} \quad (4.4)$$

where columns denote body areas, rows are the different viewpoints in each scenario, and the cell value determines whether or not the body area is visible in that frame (with a 1 meaning it is visible). This matrix is then applied during training and testing so the network can avoid learning or deciding whether a threat exists in body areas not seen in the current frame. For example, the most clear cases for this matrix would be the chest (body area 5) and upper back regions (body area 17), which will be completely hidden in the viewpoints where the subject is facing away from [Fig. 4.3(i)] and to [Fig. 4.3(a)] the scanner, respectively. Indexing M to search for these cases (column 5 with row 9, and column 17 with row 1), it can be seen that the values found are 0, indicating that the body region cannot be fully seen in those viewpoints. This can be further verified with the sample images in Fig. 4.8.

Figure 4.8: Examples from a single scenario of the DHS dataset where certain body areas are not visible in some viewpoints. (a) Shows viewpoint number 9, where body area 5 (the chest) is not visible, and (b) is the complimentary case; viewpoint 1 where body area 17 (the upper back) is not seen.

The probability of a threat existing in each body area of the labeled images in the dataset can be seen in Fig. 4.9. From this figure it can be concluded that the dataset is severely unbalanced, with a $\sim 90\%$ probability of not encountering a threat in each zone. Such a skew is heavily detrimental in deep learning, because it can cause the network to overfit to the more probable data. To avoid this effect a technique called oversampling can be used [268]. In our implementation, the oversampling consists of passing each scenario presenting a threat repeated times though the network during training to compensate the lack of more images. To achieve a more balanced distribution of scenarios with and without threats, each image with a threat is passed 9 times (to counter the $\sim 90\%$ probability).

Figure 4.9: Probability of threat existence in each body zone for the labeled images in the DHS dataset.

Because the network now encounters the same image 9 times during training, overfitting to these specific images can now occur. To solve this, during the training portion of the development, the images input into the network are randomly treated with several effects to increase image variation. These effects include added noise, slight translations and zooms, flips, and rotations. This pre-processing is a very common technique in deep learning, called

data augmentation [269], used to avoid overfitting the network on reduced training data⁴. In addition, input images during training are randomly selected and shuffled to avoid overfitting to a specific input sequence.

Finally, to improve initial learning results with the provided dataset some image processing steps have been applied to the original DHS images:

- **Colormap conversion.** Regarding the colormap of the images, only the intensity or brightness is of interest. Thus images have been converted to grayscale for two reasons. First, an equalization of the RGB channels is achieved. In addition, the network is simplified by only having to account for information from a single channel in each image.
- **Logarithmic scale.** Pixel intensity was then also converted to a logarithmic scale, in order to switch from volts to the more-traditional dB representation of returned power used, for example, in the 300 GHz and 122 GHz imaging radars. This step highlights better the contrast between low-power areas (such as the subject's body) with highly-specular zones (like metallic threats), facilitating the network's detection of oriented threats.
- **Image resizing.** In order to speed up network training and testing, images were downsampled from their original 960x760 size to 227x227 pixels (a value more common in state-of-the-art CNNs). The new symmetric size has the added benefit of allowing more efficient allocation in the GPU's memory. Together with the grayscale conversion this resizing lowers each .jpeg file size from ~150 KB to just 6 KB.

As an example, Fig. 4.10 shows a comparison between an image from the raw dataset and the same image after the processing steps listed above. Notice that the use of grayscale eliminates the red color of the returned power glints, providing a smoother representation. Additionally, after applying the logarithmic scale, the background is thresholded out, which provides better contrast for the subject's contour.

Figure 4.10: Sample comparison between an original (a) and processed (b) image from one of the DHS dataset's scenarios.

During the network's development the labeled images from the DHS dataset were split into training and validation subsets, while the unlabelled images were only used for testing. The validation subset was composed of 100 images (a 1-to-12.47 ratio), and was cycled through all labeled scenarios in subsequent trainings of the network.

To compute the predicted threat probability (p_{im}) in a specific region (i) for one of the scanning scenarios (m), the model averages the softmax outputs of each viewpoint image for

⁴As in the DHS dataset, which contains few images compared to state-of-the-art machine learning datasets.

the CNN in charge of that body area and with visibility in that region. This operation can be summarized by

$$\hat{p}_{im} = \frac{\sum_{k=1}^{16} (v \cdot \hat{p}_{ik})}{N_{visible}} \quad (4.5)$$

where $N_{visible}$ is the total number of viewpoints with visibility in that body area according to (4.4), $\hat{p}_{ik}$ is the predicted probability (softmax output) for body area i in viewpoint k , and v is a variable equal to 1 if k has visibility in i and 0 otherwise. For each scan scenario of the test subset, $(\hat{p}_{im})$ is computed for all 17 body regions.

4.1.3 Detection Results

The model has been trained and tested in the same CPU and TITAN X GPU used for the 300 GHz imaging radar in Section 2.2. Training all 17 networks on this machine lasts ~ 7.7 hours, while testing can be done in real-time. Fig. 4.11 presents a bar graph of the average duration of training for each individual network in the model, measured over 10 training processes, where the global average is 29 minutes (equivalent to 11 epochs) per network.

Figure 4.11: Average training duration in minutes for each network in the model.

The evolution of a network's accuracy throughout training can give insight regarding the correct selection of network parameters. As illustration, Fig. 4.12 shows the training accuracy evolution for three of the model's CNNs. As can be expected, accuracy starts at $\sim 50\%$, because the CNN performs a binary decision, and as the training progresses the network learns and increases its accuracy. For the first ~ 500 iterations the accuracy rises fast to almost $\sim 90\%$, as the CNN quickly learns the more general image features of the dataset. For higher accuracies the CNN needs to learn more specific information, which is a slower process requiring the next ~ 1000 iterations. Furthermore, the momentum and the learning rate drop factor can start slowing down updates as the training goes on. Notice that the curve is not smooth and presents noise. In extreme cases, this could indicate that the learning rate or batch size are set too high, but for the developed CNNs it can be seen that noise variations are reduced as the network learns, and are almost nonexistent by the time the network reaches its goal. It can also be noticed that the accuracies stop before reaching a constant 100% value, which can be naively thought of as the goal. As mentioned in the previous subsection, this is because an early stopping technique (conditioned to a 99% accuracy average) is used instead, to avoid overfitting the network to the specific input data.

The three networks shown present slightly different learning speeds (i.e. network 4 reaches higher accuracies faster than the others), which can indicate better input data or better training parameters, and which results in their duration also being different (i.e. because network 1 is the slowest learner, its duration until reaching the early stop is longer). The different learning speeds of these three specific networks can likely be explained by their corresponding body areas. According to the matrix in (4.4), network 1 (which looks at the right shoulder area) takes into account 15 viewpoints, one more than networks 4 and 14. This increases the variability of the input data and the number of images to be passed through the network, slowing down its learning. Comparing between networks 4 and 14, the difference is likely due to the physical characteristics of their body regions. Network 4 looks at the right forearm area, which is a thinner section compared to network 14's right calf, and thus, threats are more likely to stand out and be easily detected (speeding up learning).

Figure 4.12: Evolution of training accuracy for three of the model's CNNs.

The most useful tool to quantify a CNN's accuracy is its confusion matrix. This chart represents the correct and erroneous classifications among all classes, giving insight into the number of misclassifications taking place and to which wrong class they are being assigned. Fig. 4.13 presents a diagram of the information included in the confusion matrix of a 2-class problem with two generic classes (*Positive* and *Negative*). The true positives (TP) are the positive samples correctly classified as positives, and the true negatives (TN) are the negative samples correctly classified as negatives. Meanwhile, the false positives (FP) are the negative samples incorrectly classified as positives, and the false negatives (FN) are the positive samples incorrectly classified as negatives. The confusion matrix sometimes includes an additional row and column. For the additional row, the information includes the recall (TPR), false negative rate (FNR), specificity (SPC), and false positive rate (FPR), where

$$TPR = \frac{TP}{TP + FN} \quad (4.6)$$

$$FNR = \frac{FN}{TP + FN} \quad (4.7)$$

$$SPC = \frac{TN}{TN + FP} \quad (4.8)$$

$$FPR = \frac{FP}{TN + FP} \quad (4.9)$$

Predicted Class	Positive	True Positive (TP)	False Positive (FP)	Precision and False Discovery Rate
	Negative	False Negative (FN)	True Negative (TN)	
Recall and False Negative Rate				OA
		Positive	Negative	
Actual Class				

Figure 4.13: Diagram of a generic confusion matrix for a 2-class classifier.

and for the additional column, the information includes the precision (PPV), false discovery rate (FDR), negative predictive value (NPV), and false omission rate (FOR), where

$$PPV = \frac{TP}{TP + FP} \quad (4.10)$$

$$FDR = \frac{FP}{TP + FP} \quad (4.11)$$

$$NPV = \frac{TN}{TN + FN} \quad (4.12)$$

$$FOR = \frac{FN}{TN + FN} \quad (4.13)$$

Within these ratios, recall and precision provide important benchmarks. For a given actual class, the recall measures the rate at which the classifier correctly predicts its samples. Meanwhile, for a given predicted class, the precision measures how often that prediction is correct. The false negative rate and the false discovery rate are their corresponding opposite rates (the rate at which the classifier incorrectly predicts a class, and how often class predictions are incorrect, respectively). For any classifier, it is desired for recall and precision to have high values. The single cell in the lower right corner of Fig. 4.13 indicates the overall accuracy (OA) of the classifier, which can be calculated from

$$OA = \frac{TP + TN}{TP + TN + FP + FN} \quad (4.14)$$

In our implementation only two classes are available (“Si” when a threat present and “No” otherwise). The matrices for all 17 CNNs may be seen in Fig. 4.14, where the diagram from Fig. 4.13 is followed⁵. As can be seen, all CNNs present overall accuracies greater than 90%, and lower than 100% due to the early stopping condition. Furthermore, misclassifications occur in at most 1% of the image cases for any class in all the CNNs. Finally, looking at individual class percentages, it can be noticed that there is still some dataset imbalance (i.e. there isn’t the same number of images in each class) despite the oversampling technique used.

⁵Diagonal cells are correctly classified images and off-diagonal cells are misclassified cases, the lowest row and the rightmost column include the rates from equations (4.6)-(4.13), and the lower right corner provides the overall accuracy from equation (4.14).

Figure 4.14: Confusion matrices for all 17 CNNs in the developed model, in sequential order, left to right and top to bottom. The vertical axis in each chart represents the predicted class, while the horizontal axis represents the actual class. Matrix layout follows the diagram in Fig. 4.13.

Analyzing misclassified cases, it can be seen that some of the accuracy errors are due to peculiarities of the subject and scenario. For instance, subjects with abnormal joint reflections in knees and elbows caused false alarms in the corresponding body areas.

A similar analysis can be performed for the evolution of each network's loss during training. The loss results for the same three CNNs in Fig. 4.12 can be seen in Fig. 4.15. As expected, the speed at which the loss decreases is correlated with the speed at which the network's accuracy increases. Due to this, network 4 is once again the CNN which reaches better results faster.

Figure 4.15: Evolution of training loss for three of the model's CNNs.

The loss achieved by each of the model's networks when they have finished their training is summarized in Fig. 4.16. The training results in the graph are much lower than the model's 0.16795 testing loss obtained in the Kaggle challenge. An increase in testing loss compared to training loss is common and to be expected, because the test dataset can include scenarios and features unseen during training, but such a large different could indicate overfitting to the training data.

Figure 4.16: Training loss for each CNN in the developed model.

Comparing each network's losses among themselves, it can be noticed that the values

present significant variance. This suggests that specific networks have performed better than others. In this regard, network 5 (in charge of the upper chest area) presents the best loss. In similar vein, individual losses of 0.1221 and 0.0088 for the upper chest and upper back body areas have also been reported by other authors for the same dataset also using AlexNet [270]. These results are worse than our equivalents in Fig. 4.16 (networks 5 and 17, respectively), indicating that the modifications performed to our AlexNet implementation have improved performance. Different authors have reported accuracies of 98.28% with 0.091319 loss also using AlexNet, but on a subset of 200 scenarios from the DHS dataset [271], which suggests that better performance could be achieved with limited image scenarios. Finally, in another approach on the DHS dataset [272], individual losses between 0.1696 and 0.3357 were registered for the different body areas using a VVG-16 CNN [273], which are higher than our achieved scores, validating our choice of an AlexNet implementation.

In regard to the training, much of the effort dedicated to the development of the model has been focused on the optimum selection of network parameters. Weight initialization and training parameters (such as learning and dropout rates) are crucial to the learning process, and produce very different results with slight modifications. Although the intuitive nature of these parameters allows easy selection of an adequate starting value, their fine tuning is still commonly done with trial-and-error approaches. Time-wise, this is extremely costly due to the extended duration of the network's training stage. Nonetheless, the correct setting of these parameters produces significant improvements. As illustration example, the following timeline in loss [S , according to (4.1)] decreases was registered:

- An initial lightly-modified AlexNet implementation reached a loss of $S = 0.69527$.
- The use of data augmentation strategies such as image translation, random crops, and rotations, produced a reduction in loss from $S = 0.69527$ to $S = 0.3817$.
- In addition, training the networks using only frames where the threat could be visible according to the matrix in (4.4) reduced the loss to $S = 0.20392$.
- Afterwards, image preprocessing, such as grayscale conversion, further improved the score to $S = 0.19128$.
- Randomly initializing the weights of the later AlexNet layers, instead of using the ImageNet-trained values, lowered the loss to $S = 0.18113$.
- Finally, the introduction of oversampling to balance the dataset's distribution and the fine-tuning of training hyperparameters produced a loss of $S = 0.16795$.

Visualizing the final weights of the network as images can provide useful information regarding the network's learning. The weights for each filter in a given layer can present visual information about what image features the filter is focusing on. For instance, filters that provide smooth shapes without noisy patterns indicate that the learning process has converged naturally without abrupt steps (which would create sharp features in the filter) and is likely to be in a stable state. These visualizations can help decide whether the training stage has been sufficiently long or if overfitting has occurred. The visualization of a layer's weights is more intuitive earlier in the network (in the initial layers), where the filters have converged to detect very general and easily-distinguishable image features (such as edges and blobs). If done in deeper layers of the network, the represented shapes are likely to be more abstract and less easy to interpret, since the later layers have built on the information from the initial stages. As illustration, Fig. 4.17 presents representations of the weights of the 96 filters belonging to the first convolutional layer of each of the 17 CNNs in the developed model.

Figure 4.17: Weight representation for the first convolutional layer of all 17 CNNs in the developed model, in sequential order, left to right and top to bottom.

As can be seen from the results, the filters have converged to learn nondescript shapes, and their value is approximately the same throughout all CNNs. This is because, although each CNN is in charge of a specific body region, the first convolutional layer is always only able to learn primitive features. Deeper labels in the network build and combine previously learnt features into more specific filters. Notice that the filters can be broadly classified into two categories: colored filters with rounded and smooth contours, and grayscale filters with sharper lines and shapes. This is an inherent effect of using an AlexNet architecture, which divides its processing into two streams (as seen in Fig. 4.7), where one stream ends up developing high-frequency grayscale image features while the other develops low-frequency color features [263].

A layer's activation maps are another useful tool to evaluate the effectiveness and efficiency of the network's training. These maps are the outputs of a layer in the network after an image is passed through it, and show whether or not a neuron (filter) has been activated by the information in that specific image input. If the output pixels in the body region assigned to each network are never activated through the full dataset this can indicate the existence of dead paths in the CNN. To test whether or not the model is learning from the appropriate image features, an analysis for network 5 (in charge of the upper chest area) is performed. In order to test this CNN, a processed image (shown in Fig. 4.18) from a scenario in the DHS dataset where the subject is carrying a threat in body area 5 is used. As can be seen, the subject is carrying a large oval threat in his chest, where the selected network should have been trained to look.

Figure 4.18: Sample processed image from the DHS dataset showing a person with an oval threat in the upper chest area (body region 5). Threat boxed in blue, only for this display

The first test is the activation map of the network's first convolutional layer. The layer has 96 channels⁶, each of which corresponds to a 2-D image showing the activations of said channel. When plotted in a grid format, the resulting activation maps for the image in Fig. 4.18, may be seen in Fig. 4.19. Brighter pixels represent strong positive activations and blacker pixels represent strong negative activations, while gray pixels represent insignificant activation. As can be seen from the grid results, the first layer is mostly in charge of detecting general image features such as edges or contours (which correlates with the generic shapes seen in Fig. 4.17). In deeper layers activations should become even more sparse, as the layers are focused on detecting specific features, building up on the learning done in the previous layers. The channel with strongest activations is compared to the original image input that produced it (Fig. 4.18) in Fig. 4.20, where it can be noticed how clearly the filter has profiled the subject's contour

⁶The fourth dimension of a layer's size, which indicates the number of filters (neurons) in each layer.

Figure 4.19: Activation maps for network 5's first convolutional layer when input with the image in Fig. 4.18.

Figure 4.20: Comparison of the image from Fig. 4.18 and the maximum activation map from Fig. 4.19.

If the same type of representation is repeated for a deeper CNN layer, for example the fifth and last convolutional layer, the resulting grid can be seen in Fig. 4.21. It is obvious from the results, comparing to the grid in Fig. 4.19, that the image features learnt in this layer are not as intuitive as in the first convolutional layer. Nonetheless, the channel with maximum activations in this case also presents an interesting feature. A comparison between this channel and the input image may be seen in Fig. 4.22. It can be noticed that the channel presents a bright circular spot indicating strong positive activation, but also some gray and black areas. This representation might not seem very informative, and since after the convolutional layer a ReLu function layer is used [which eliminates negative activations, through the formula in (4.3)], it can be more representative to display the channel's output from the following ReLu layer rather than the convolutional layer. If this is done, the result is as shown in Fig. 4.23(a), where the black and gray activations (which represent negative and null values, respectively) are eliminated by the ReLu function and only the strong positive white spot is left. If this activation map is converted to a red color and overlaid on the input image, as in Fig. 4.23(b), it can be seen that the channel's strong activation is due to pixels in the subject's chest area, which is the exact region where the threat is placed, and where the CNN should be trained to focus. This result suggests that the network is indeed learning.

To further confirm this learning, it is useful to test a negative case. For this task, another image is pulled from the DHS dataset, shown in Fig. 4.24, displaying a subject with the same body type as Fig. 4.18, but with no threat in his chest area. The activation map of the same channel used in Fig. 4.23, for the new input image, can be seen in Fig. 4.24. This time, no strong activations are present, indicating that the CNN has learned to detect that no threat is located in the chest area.

Figure 4.21: Activation maps for network 5's fifth convolutional layer when input with the image in Fig. 4.18.

Figure 4.22: Comparison of the image from Fig. 4.18 and the maximum activation map from Fig. 4.21.

Figure 4.23: (a) Comparison of the image from Fig. 4.18 and the maximum activation map from the network's ReLu layer after the fifth convolutional layer, (b) along with an overlay of the two after coloring the activation map in red.

Figure 4.24: Comparison of a processed image from the DHS dataset showing a person with no threat in the upper chest area (body region 5) and the maximum activation map from the network's ReLu layer after the fifth convolutional layer.

4.2 ATD in Passive Infrared Images

The ATD model developed for active images in the previous section can be also adapted to passive imaging results. Such a model would be useful for the passive sensor included in the hybrid imager described in Chapter 3. As of the writing of this thesis, the lack of real-time capabilities of the hybrid system creates a significant bottleneck to build a large dataset with enough variation in imaging scenarios. As mentioned in the previous section, traditional state-of-the-art datasets can include millions of images, which is not a feasible goal for the current hybrid imager and thesis timeline. For this reason, and similarly to the active dataset in Section 4.1, the source for the passive images in this section has not been one of the systems characterized in this thesis.

4.2.1 Image Dataset and Network Architecture

The passive dataset has been created using the IR camera presented in Chapter 3. It is an E60 model from FLIR Systems, which provides IR images at wavelengths between $7.5 \mu\text{m}$ and $13 \mu\text{m}$ with a thermal sensitivity of 0.05°C [274]. Its images have 320×240 dimensions and represent temperatures from -20°C to 650°C at a 30 fps frame rate. The device has a built-in display, but can also store images in memory or send them via USB or WiFi.

Although several public IR datasets exist (some of them are referenced in Table 4.2), none were found to be specifically apt for standoff concealed threat detection. Most datasets are focused on applications related to pedestrian and vehicle movement, with functionalities such as detection, counting, and tracking. This is because, in contrast to conventional cameras, IR devices can function equally in most daytime and nighttime conditions, which allows 24-hour operability.

Table 4.2: IR imaging datasets

Reference	Application	Dataset Size
Our work	Concealed threat detection	~ 170000 images
VLIRVDIF [275]	Video and IR registration	24 video sequences*
OTCBVS [276]	Various [†]	Depends on dataset [‡]
Kristoffersen et al. [277]	Pedestrian counting	3 video sequences of 5 min each
LTIR [278]	Object tracking	20 video sequences*
BU-TIV [280]	Object tracking	~ 23000 images
CATS [281]	Video and IR registration	1372 images
CAMEL [282]	Object tracking	~ 43000 images

*Roughly 2000 images per sequence.

[†]Such as face analysis, pedestrian tracking, or vehicle detection.

[‡]Different datasets within the collection have different sizes (as little as ~ 4000 images to as large as 20 video sequences).

*Average length of 563 frames per sequence.

For example, the images in VLIRVDIF [275] include many pedestrian targets, but only a handful of the sequences present concealed objects. Within the OTCBVS collection [276], the most appropriate datasets for our application would be the OSU Thermal Pedestrian Dataset, the OSU Color-Thermal Database, the Terravic Motion IR Database, and the Terravic Weapon IR Database. The first three are focused on person-detection and tracking, not placing any importance on concealed weapons, so they could at most be used for negative cases (i.e. subjects with no threats) in our application. The latter one is also not apt because it includes only 5 video sequences, where the weapon is not hidden, and the focus of the dataset is to detect the weapon's discharge while firing. The dataset developed by

Kristoffersen [277] describes an interesting stereo IR imaging set-up, but it is applied only to counting and tracking pedestrians passing through an area. The LTIR dataset [278], used in the VOT-TIR challenge since 2015 [279], contains images of several targets (such as animals, humans, and drones), but is also focused solely on tracking them. Similarly, the BU-TIV dataset includes varied targets for tracking [280], but no concealed weapon scenarios. Much the same as the VLIRVDIF collection, the CATS and CAMEL datasets include matching IR and conventional camera outputs [281, 282], but no threats are present in these cases. Because no apt public dataset was found, a custom-made collection of images focused on threat detection scenarios was required.

Our dataset's images are generated in real-time through the camera's video capabilities. For each imaging scenario, a 4 minute video is captured, providing ~ 7200 useful frames. Extracting individual images as frames from a video allows generating large datasets quickly, at the cost of introducing images that are very correlated (which can produce overfitting during learning). To counteract correlation, variations in several areas of the scenarios are introduced, along with the data augmentation techniques used in Section 4.1. The dataset includes 7 different threat detection classes, corresponding to 6 different threats (shown in Fig. 4.25) plus a no-threat option.

Figure 4.25: Threats included in the IR dataset, in order from left to right: ice pick, 9 mm handgun, ceramic knife, shrapnel explosive (plasticine+bolts), water bottle, IED (capacitors).

Two male subjects with different body types are imaged in the dataset. Clothing varies in two options: light cotton T-shirts or heavy wool sweatshirts. Threat temperature, which defines contrast in IR images, is also varied by either cooling or warming threats before imaging. The threats can be placed in various body areas, including torso, arms, and pockets, among others. Videos are captured in an air-conditioned room with no clutter in the scene.

The subject starts each video sequence facing the camera, which is placed in a fixed tripod ~ 3 m away. This set-up is chosen to simulate the usual viewpoint in security checkpoint, but the subject is not asked to maintain a specific pose throughout the sequence. Instead, the subject can move his arms, head, and torso freely during 1 minute while facing the camera. After every minute, the subject is simply asked to perform a $\sim 90^\circ$ clockwise rotation, which after 4 minutes allows capturing a full 360° scan. Fig. 4.26 provides samples frames from one of these sequences. This rotational set-up introduces some difficulties. Three frames in Fig. 4.26 clearly show the bottle as a dark blue shape in the warm torso, so they (and all others where the posture is maintained) are labeled as threat scenario. Meanwhile, Fig. 4.26(c), does not show the threat, so this image is labeled as a no threat scenario. The problem is deciding at which specific frame during the rotations from Fig. 4.26(b) to Fig. 4.26(c), and from Fig. 4.26(c) to Fig. 4.26(d), the threat fully appears or disappears from the image. This is illustrated by the samples in Fig. 4.27, which are produced just 1 second apart, but because the subject is moving the threat is no longer fully visible in the second image. If it is labeled as a threat scenario, the CNN would be forced to classify from incomplete shapes, which would lower its accuracy. To avoid this, only frames that fully show the object are manually labeled as a threat scenario. Since frames are taken at a 30 fps rate choosing the exact image at which the threat is fully shown can be problematic, so some error is expected.

Figure 4.26: Sample frames from an imaging scenario in the IR dataset where a male subject hides a cool water bottle underneath a T-shirt. The frames show the four $\sim 90^\circ$ viewpoints the subject is asked to maintain during one minute each. Colorbar: Temperature in degrees.

Figure 4.27: Sample frames from the scenario in Fig. 4.26 taken just 1 second apart. Frame (a) is labeled as a threat scenario because it displays the full bottle, while frame (b) is labeled as a no threat scenario because it only shows part of the object. Colorbar: Temperature in degrees.

In contrast to the active dataset in Section 4.1, no preprocessing is done on the images before inputting them into the network. The only modification performed is a crop of the supplementary information the camera adds into the image (such as colorbar and time), done to avoid introducing unnecessary shapes that the CNN would be learning from. Most public IR datasets do not include these data either [275, 276]. A comparison between a raw frame and a cropped result is presented in Fig. 4.28.

Figure 4.28: (a) Raw IR frame and (b) cropped result. In the cropped example all unnecessary shapes in the image have been erased. Colorbar: Temperature in degrees.

The full dataset contains 165278 images, distributed into the 7 threat classes as shown in Fig. 4.29. The scenario where no threat is present ("None") includes 53329 images, while all other threat classes range between 6995 and 8201 images. This large difference is due to the fact that no threat images can be extracted from the frames in video sequences where the threat is simply not visible, but they have also been generated from full video sequences where no threat has been used. Oversampling during training, such as in Section 4.1, has been used to compensate this unbalance.

Figure 4.29: Distribution of images from the IR dataset into the 7 threat classes considered.

As previously mentioned, image correlation due to generating dataset entries as video frames could be a problem. To attempt to quantify the image diversity of the full dataset, the approach proposed for ImageNet in [237] has been followed. In this method, diversity is quantified as the file size of the lossless JPG image corresponding to the average of all dataset frames. The size of the file is a measure of the amount of information in the average

image, which in the ideal case (maximum diversity) will simply be a blurry and fully-gray image. For our dataset, the file size has been 2 KB, which approximately doubles the sizes seen for ImageNet and other conventional datasets [237]. This was to be expected, as these other datasets are created with individual images (not video sequences) coming from multiple external sources (not from a single imaging device, scene, or date). Fig. 4.30 presents the resulting average image of the full dataset, where it is apparent that the dataset presents high correlation from the facts that a human body is still somewhat discernible and the background is approximately constant.

Figure 4.30: Average image of the IR dataset generated for image diversity calculation.

The CNN used to detect threats in this dataset follows a similar architecture to the one in Section 4.1. The network is based around an AlexNet implementation, pre-trained on ImageNet⁷, and with slight modifications to the last layers to allow transfer learning (randomized weights as starting point for the last three convolutional layers). The size of the last layer has been set to seven outputs in order to distinguish between all possible threat scenarios. In contrast to the design for the active dataset, a single CNN is used, receiving images from all viewpoints (instead of using an individual CNN for each of the subject's positions).

4.2.2 Detection Results

The images from the dataset are split in a 4:1 ratio for training and validation, respectively. As for the active images, the CNN is set for early stopping during training when it reaches a 99% accuracy goal to avoid overfitting. Training hyperparameters, such as momentum, initial learning rate, and batch size have been set to the same values used in Section 4.1.

The model has first been tested including only the simplest scenarios: images of a single subject hiding a cooled threat between his torso and T-shirt. This narrows down the training images to ~ 78000 frames, for which the CNN requires an average of 10.5 minutes of training on an NVIDIA TITAN X GPU. Fig. 4.31 shows the training loss and accuracy for the CNN and these inputs.

The model achieves a 99.8% accuracy and a 0.033 loss in training. For validation, the number of images is ~ 19500 , and the model rises to a 99.92% accuracy. These numbers

⁷Pre-made CNNs, with transfer learning from ImageNet, have been demonstrated to have good performance on IR images [283]

indicate great performance, specially considering the added difficulties that image classification algorithms seem to encounter with man-made and small objects [284]. For this model the confusion matrices have also been generated. Since the problem now requires multi-class classification, the confusion matrix layout from Fig. 4.13 must be modified, as seen in Fig. 4.32.

Figure 4.31: Training (a) accuracy and (b) loss for the CNN when inputting images of a single subject hiding cooled threats placed in his torso underneath a T-shirt.

Figure 4.32: Diagram of a generic confusion matrix for a N -class classifier focusing on class c_k .

The confusion matrices for training and validation in this model are shown in Fig. 4.33. Notice the accuracy for any class is never lower than 99% and less than 60 frames at most are misclassified within a single class. If the specific errors of those cases are analyzed, it can be found that most misclassifications are due to frames captured during the subject’s rotation between back-facing and side-facing positions. As illustrated with Fig. 4.27, the threat in these positions can sometimes be partially obscured or deformed, which leads the CNN to classify it as a different object or even a no threat scenario.

Output Class	Bottle	6549 8.4%	0 0.0%	0 0.0%	6 0.0%	7 0.0%	0 0.0%	28 0.0%	99.4% 0.6%
	Gun	1 0.0%	5800 7.4%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	100.0% 0.0%
	IED	1 0.0%	1 0.0%	5582 7.2%	0 0.0%	16 0.0%	0 0.0%	23 0.0%	99.3% 0.7%
	Knife	0 0.0%	3 0.0%	0 0.0%	5844 7.5%	3 0.0%	0 0.0%	0 0.0%	99.9% 0.1%
	None	9 0.0%	6 0.0%	2 0.0%	12 0.0%	42629 54.6%	4 0.0%	4 0.0%	99.9% 0.1%
	Pick	0 0.0%	0 0.0%	0 0.0%	0 0.0%	5 0.0%	5698 7.3%	0 0.0%	99.9% 0.1%
	Shrapnel	0 0.0%	0 0.0%	11 0.0%	0 0.0%	2 0.0%	0 0.0%	5776 7.4%	99.8% 0.2%
		99.8% 0.2%	99.8% 0.2%	99.8% 0.2%	99.7% 0.3%	99.9% 0.1%	99.9% 0.1%	99.1% 0.9%	99.8% 0.2%
		Target Class							
		Bottle	Gun	IED	Knife	None	Pick	Shrapnel	

(a)

Output Class	Bottle	1640 8.4%	0 0.0%	0 0.0%	0 0.0%	2 0.0%	0 0.0%	3 0.0%	99.7% 0.3%
	Gun	0 0.0%	1451 7.4%	0 0.0%	0 0.0%	1 0.0%	0 0.0%	0 0.0%	99.9% 0.1%
	IED	0 0.0%	0 0.0%	1399 7.2%	0 0.0%	3 0.0%	0 0.0%	0 0.0%	99.8% 0.2%
	Knife	0 0.0%	1 0.0%	0 0.0%	1465 7.5%	4 0.0%	0 0.0%	0 0.0%	99.7% 0.3%
	None	0 0.0%	0 0.0%	0 0.0%	1 0.0%	10654 54.6%	0 0.0%	0 0.0%	100.0% 0.0%
	Pick	0 0.0%	0 0.0%	0 0.0%	0 0.0%	1 0.0%	1425 7.3%	0 0.0%	99.9% 0.1%
	Shrapnel	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	1455 7.5%	100% 0.0%
		100% 0.0%	99.9% 0.1%	100% 0.0%	99.9% 0.1%	99.9% 0.1%	100% 0.0%	99.8% 0.2%	99.9% 0.1%
		Target Class							
		Bottle	Gun	IED	Knife	None	Pick	Shrapnel	

(b)

Figure 4.33: Confusion matrices for (a) training and (b) validation for the CNN when inputting images of a single subject hiding cooled threats placed in his torso underneath a T-shirt. The vertical axis in each chart represents the predicted class, while the horizontal axis represents the actual class. Matrix layout follows the diagram in Fig. 4.32.

For the next training stage, additional variation was included by introducing images from a different male subject. The subject still hides cooled threats in his torso underneath a T-shirt, but his height, body type, and clothing present some differences. Fig. 4.34 compares an imaging scenario where the threat is a handgun for the two male subjects.

Figure 4.34: Comparison between the two male subjects considered in the dataset for the same threat scenario (a concealed handgun).

When the trained model from Fig. 4.31 is tested with ~ 77000 images from the new male subject, the CNN's overall accuracy decreases to 88.2%. This can be seen in the testing confusion matrix presented in Fig. 4.35(a). The CNN performs specially worse on the ice pick ("Pick") and knife ("Knife") classes, often confusing them with one another, as well as with the no threat ("None") scenario. This is due to the similarity in shape between these two threats, which in addition have their contours distorted when the clothing over them starts becoming cold through contact. As for their misclassification as a no threat scenario, because they are much thinner than the remaining targets, depending on the subject's body posture their temperature contrast can be reduced in size and thus undetected. To solve this loss in performance, the CNN can be retrained with the combined set of images generated from both subjects. When receiving inputs from both subjects during training, the model is forced to learn less dataset-specific image features, which provides more generalized learning. The results of following this training option can be seen in Fig. 4.35(b), where the model once again obtains 99.2% accuracy.

Ultimately, any additional variation included to the images, such as the difference in subjects seen before, will be less detrimental to the CNN's performance if the model is retrained including new images. On the other hand, for real-life implementation the model is required to have learned enough general features to be able to detect threats in image scenarios it has not seen before. For this purpose, the dataset needs to grow and account for as much variations as possible. In addition to introducing another subject, new image sequences were also captured with uncooled threats, hidden beneath thick sweatshirts, and placed in random body areas. Fig. 4.36 shows examples from these three new image variations.

If the model from Fig. 4.35(b) is not retrained and is instead directly tested with ~ 86000 images which include these new variations, its resulting performance may be seen in Fig. 4.37(a). In this case, a large decrease in overall accuracy to just 74% has been produced. More importantly, only the no threat class achieves a probability higher than 41%, with the "Bottle", "Gun", and "Knife" classes obtaining 0% detection, which seems to indicate that the CNN is biased to the no threat scenario. From these results we can extract some considerations:

Output Class	Bottle	6735 7.7%	142 0.2%	39 0.0%	736 0.8%	124 0.1%	84 0.1%	93 0.1%	84.7% 15.3%
	Gun	0 0.0%	5222 6.0%	34 0.0%	10 0.0%	4 0.0%	0 0.0%	0 0.0%	99.1% 0.9%
	IED	0 0.0%	483 0.6%	6941 8.0%	0 0.0%	45 0.1%	0 0.0%	465 0.5%	87.5% 12.5%
	Knife	0 0.0%	2 0.0%	0 0.0%	4121 4.7%	0 0.0%	2913 3.3%	0 0.0%	58.6% 41.4%
	None	100 0.1%	49 0.1%	25 0.0%	1071 1.2%	44449 51.1%	1254 1.4%	30 0.0%	94.6% 5.4%
	Pick	0 0.0%	6 0.0%	0 0.0%	1087 1.2%	4 0.0%	2781 3.2%	7 0.0%	71.6% 28.4%
	Shrapnel	111 0.1%	1205 1.4%	56 0.1%	0 0.0%	48 0.1%	0 0.0%	6520 7.5%	82.1% 17.9%
		97.0% 3.0%	73.5% 26.5%	97.8% 2.2%	58.7% 41.3%	99.5% 0.5%	39.5% 60.5%	91.6% 8.4%	88.2% 11.8%
		Target Class							
		(a)							

Output Class	Bottle	5859 8.1%	6 0.0%	2 0.0%	0 0.0%	73 0.1%	0 0.0%	12 0.0%	98.4% 1.6%
	Gun	4 0.0%	5688 7.9%	0 0.0%	1 0.0%	17 0.0%	1 0.0%	0 0.0%	99.6% 0.4%
	IED	0 0.0%	6 0.0%	5686 7.9%	0 0.0%	9 0.0%	0 0.0%	53 0.1%	98.8% 1.2%
	Knife	2 0.0%	11 0.0%	0 0.0%	5731 8.0%	66 0.1%	4 0.0%	0 0.0%	98.6% 1.4%
	None	0 0.0%	10 0.0%	0 0.0%	1 0.0%	37117 51.5%	9 0.0%	3 0.0%	99.9% 0.1%
	Pick	0 0.0%	5 0.0%	0 0.0%	11 0.0%	256 0.4%	5659 7.9%	0 0.0%	95.4% 4.6%
	Shrapnel	2 0.0%	0 0.0%	14 0.0%	0 0.0%	30 0.0%	0 0.0%	5656 7.9%	99.2% 0.8%
		99.9% 0.1%	99.3% 0.7%	99.7% 0.3%	99.8% 0.2%	98.8% 1.2%	99.8% 0.2%	98.8% 1.2%	99.2% 0.8%
		Target Class							
		(b)							

Figure 4.35: Confusion matrices during testing for the CNN when inputting images of two male subjects when the training images come from (a) a single subject or (b) both subjects. The vertical axis in each chart represents the predicted class, while the horizontal axis represents the actual class. Matrix layout follows the diagram in Fig. 4.32.

Figure 4.36: Sample images of additional scenario variations considered in the full IR dataset, such as (a) uncooled threats, (b) thicker clothing, or (c) random threat placements. In all images the subject is hiding the same handgun.

- To detect threats, the model had learned to rely more on their color component (i.e. their temperature contrast with their surroundings) rather than their shape. This is the reason why including uncooled or heavily-covered threats (which present approximately the same color as the body or clothing) are no longer being detected, despite the fact that they can still be visually detected [as illustrated by Fig. 4.36(a)].
- The location of the threat has also negatively affected performance. This is because the CNN had been trained with images where the threat was always located in the torso's center and with the same rotation. This implies that the model had not learned to recognize the threats' general shapes, but rather specific placements.

Once again, to solve this obstacle the model is retrained with the full IR dataset, which includes images from all variations. Training in this case requires an average of 33.43 minutes. If the model is now tested with ~ 72000 images from the full dataset, its performance may be seen in Fig. 4.37(b), where the overall accuracy is 99.4% (on par with the benchmarks obtained with the individual subsets). This means that a CNN capable of detecting 7 different threat classes with varying temperature, placed in random body locations with two different subjects and types of concealing clothing, has been developed.

Output Class	Bottle	26 0.0%	0 0.0%	18 0.0%	0 0.0%	8 0.0%	0 0.0%	4606 5.3%	0.6% 99.4%
	Gun	0 0.0%	0 0.0%	24 0.0%	0 0.0%	24 0.0%	0 0.0%	0 0.0%	0.0% 100%
	IED	0 0.0%	2 0.0%	122 0.1%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	98.4% 1.6%
	Knife	0 0.0%	0 0.0%	0 0.0%	0 0.0%	1 0.0%	0 0.0%	1 0.0%	0.0% 100%
	None	1320 1.5%	4974 5.7%	81 0.1%	4461 5.2%	60350 69.7%	2541 2.9%	54 0.1%	81.8% 18.2%
	Pick	0 0.0%	0 0.0%	0 0.0%	2 0.0%	89 0.1%	1158 1.3%	1 0.0%	92.6% 7.4%
	Shrapnel	3998 4.6%	293 0.3%	56 0.1%	0 0.0%	2 0.0%	1 0.0%	2405 2.8%	35.6% 64.4%
		0.5% 99.5%	0.0% 100%	40.5% 59.5%	0.0% 100%	99.8% 0.2%	31.3% 68.7%	34.0% 66.0%	74.0% 26.0%
		Target Class							
		Bottle	Gun	IED	Knife	None	Pick	Shrapnel	

(a)

Output Class	Bottle	5541 7.7%	1 0.0%	0 0.0%	6 0.0%	74 0.1%	0 0.0%	23 0.0%	98.2% 1.8%
	Gun	3 0.0%	5349 7.4%	0 0.0%	0 0.0%	16 0.0%	0 0.0%	0 0.0%	99.6% 0.4%
	IED	0 0.0%	2 0.0%	4335 6.0%	0 0.0%	16 0.0%	0 0.0%	17 0.0%	99.2% 0.8%
	Knife	0 0.0%	0 0.0%	0 0.0%	5187 7.2%	48 0.1%	1 0.0%	0 0.0%	99.1% 0.9%
	None	3 0.0%	2 0.0%	0 0.0%	32 0.0%	40502 56.3%	62 0.1%	2 0.0%	99.8% 0.2%
	Pick	0 0.0%	0 0.0%	0 0.0%	7 0.0%	34 0.0%	4943 6.9%	0 0.0%	99.2% 0.8%
	Shrapnel	0 0.0%	4 0.0%	8 0.0%	0 0.0%	38 0.1%	0 0.0%	5672 7.9%	99.1% 0.9%
		99.9% 0.1%	99.8% 0.2%	99.8% 0.2%	99.1% 0.9%	99.4% 0.6%	98.7% 1.3%	99.3% 0.7%	99.4% 0.6%
		Target Class							
		Bottle	Gun	IED	Knife	None	Pick	Shrapnel	

(b)

Figure 4.37: Confusion matrices during testing for the CNN when inputting images from the full IR dataset, after it has been trained with (a) a subset or (b) the full dataset. The vertical axis in each chart represents the predicted class, while the horizontal axis represents the actual class. Matrix layout follows the diagram in Fig. 4.32.

4.3 Future Work

In conclusion, two proof-of-concept ATD algorithms have been developed for active and passive images generated in security screening scenarios. Both models present excellent performance, reaching threat detection accuracies higher than 99%. The results demonstrate the feasibility of quickly adapting pre-trained CNNs designed for conventional photographic data into ATD models suited for millimeter-wave active and passive images. Nonetheless, it must be taken into account that the scenarios considered in the active and passive datasets are limited in various areas (such as number of images, threat classes, and other aspects).

With these limitations in mind, this section details some of the lessons learnt from this chapter's work and future research opportunities available to improve the ATD models presented. First, general aspects applicable to both the active and the passive models are listed. Afterwards, specific developments for the ATD model using active images in Section 4.1 are summarized. Lastly, similar improvements for the ATD model using passive images in Section 4.2 are described.

Mainly, the priority for future work should without doubt be the transfer of the developed algorithms to detect threats in the 300 GHz imaging radar from Chapter 2 and the hybrid imager from Chapter 3. Some considerations regarding this development can be made:

- Compared to the 300 GHz and hybrid systems, the images from the DHS dataset are generated using a scanner with a 360° view, at a lower operating frequency, and with different resolutions in both range and cross-range. These factors limit the certainty that the DHS-trained CNN in Section 4.1 will automatically be able to detect threats in the 300 GHz or 122 GHz active images.
- In contrast, the CNN trained in Section 4.2 has shown good performance for passive IR images, which have similar properties to the 94 GHz passive images and are generated in a comparable imaging set up (target is front-facing and at a standoff distance). In addition, the CNNs in Sections 4.1 and 4.2 follow the same core architecture (an AlexNet implementation), which in this work has been demonstrated to perform well with both active and passive images. This seems to suggest that the same CNN model could be used for the imaging results of the hybrid system.
- Although the DHS-trained or IR-trained weights will likely serve as a better starting point than an ImageNet-trained network, the CNN will still need to be tested and fine-tuned to the peculiarities of each imager. To properly do so, the first step should be the development of a relatively large dataset of images with each scanner. Given the experience with the DHS dataset, an adequate number of scenarios would be around 1000 different cases, although the higher the number the better the learning.
- Furthermore, the 300 GHz and 122 GHz active sensors can also generate range-based images (complementing the power-based results similar to the DHS dataset), which provides additional information for the network's predictions, but could require modifications in its architecture.

Aside from that main priority, both deep learning models (for active and for passive images) could benefit from the following improvements:

- Transfer learning from pre-trained networks has proven to be a successful technique to bypass limitations during learning due to a small dataset. In this regard, the weights from an ImageNet-trained AlexNet CNN provided a good starting point for the initial layers of the networks used for the active and passive images. Nonetheless, both the DHS and IR datasets have very specific image scenarios (a single standing human

subject where a manmade object is to be detected). With this in mind, pre-training the CNN on a dataset different from ImageNet, and more specific to the intended security application, could provide better performance. For this task, multiple photographic and publicly-available datasets of human targets for pose estimation [285], activity recognition [286], or medical imaging [287] could be used. Several IR datasets including images of humans are also public (as summarized in Table 4.2), and these images could even be added to the “None” threat class of the developed model to test whether the CNN has learned to detect threats generally or only in the specific scenarios of our IR dataset.

- In similar vein, despite the fact that the network selected (AlexNet) has had undeniable success in this task, as well as other object detection applications, it cannot be omitted that the network was originally designed for ImageNet-based scenarios. Because the ImageNet and threat detection datasets are considerably different, it can be expected that a custom-made network could provide better performance. A shallower network would allow removing all the ImageNet-trained layers and training with only the DHS or IR data. In addition, mm-wave images generally contain less information than photographs (for instance, regarding color), so a smaller network is again more adequate. In case a custom-designed network proves ineffective, the use of networks more modern than AlexNet, such as GoogLeNet [288], ResNet [289], or SENet [290] (all of which outperformed AlexNet in the ImageNet challenge) could provide better results.
- The threat detection application could be combined with a threat localization functionality. This task is very common in machine learning applications, and is even a sub-challenge on the ImageNet dataset (where participants detect the object in the scene and then place a bounding box around its location). Such an approach has already been reported on the DHS dataset [271] by segmenting the input images in pre-defined grids. For the IR dataset, this added functionality would require increasing the labeling information to include threat coordinates. In a real-world implementation, a threat-localizing model speeds up passenger throughput by providing the human operators with further information for a more confined manual screening.
- The use of MATLAB for development severely limited the customization and functionalities of the network. Environments more specific to deep learning, such as Keras [291], TensorFlow [292], Caffe [293], or NVIDIA’s DIGITS [294], allow in-depth tuning which will be necessary to optimize the CNN.

4.3.1 ATD in Active Images

For the ATD algorithm developed with the active images from the DHS dataset the following improvements are indicated:

- A potential source for higher accuracy could be the use of multi-view CNNs to account for all viewpoints of a single dataset scenario [295]. This kind of networks are specifically designed for shape recognition in 3-D images, which perfectly fits the threat detection task. Such an approach on the DHS dataset using a VGG-16 network reported accuracies higher than 95% [296]. For such an implementation it could be useful to integrate the other filetypes provided in the DHS dataset, which contained 3-D representations.
- The use of the subjective view matrix M from (4.4) resulted in a 46.56% decrease in loss, but this matrix was defined by looking at a single scan scenario. A more unambiguous method would be to define a matrix M_i for every scenario (i) in the dataset. This way, for all scenarios, the network only learns from the exact frames where the threat is visible.

- On a related note, the use of oversampling to balance threat probability in the dataset provided a 6.28% reduction in loss, but this was done by compensating a $\sim 90\%$ no-threat probability in each body region. However, looking at Fig. 4.9 it can be seen that there are slight differences in the exact probabilities for each body area. For instance, the probability in body area 1 (right bicep) is $\sim 12\%$ while in body area 9 (groin) it is $\sim 8\%$. Defining area-specific ratios for the oversampling could provide a more balanced dataset, and thus better learning.
- Several other minor modifications could provide more benefits. Additional image pre-processing, like histogram equalization [270], edge detection [271], or cropping body regions for their specific network [270], have been successfully tested in the DHS dataset. Changes to the network, such as a different activation function, pooling strategy, or dropout factor have already been tried during development, but could be further studied.

As reference for future improvements it can be useful to analyze the top submission in the DHS challenge [297], which achieved a loss of only $S = 0.02417$. The final design used a ResNet-50 implementation pre-trained in ImageNet. In addition to the data augmentation techniques used for our implementation, the top design included Gaussian blurring and brightness adjustments to add even more variation to the inputs. To compensate the low number of scan scenarios, the designer manually modified some of the images to translate threats from one body area to another, and thus create ~ 800 new scenarios. Instead of simply converting images to grayscale and only training the CNN with a gray channel, the winning design made use of the network's two remaining channels in a classical RGB configuration to include complementary information of each frame via its statistics (such as mean and standard deviation of neighboring pixels). This design cropped images before learning, but instead of segmenting each body area, the winners found better results by using only 4 bigger and overlapping sections (approximately covering arms, chest, waist, and feet, respectively). Finally, the final probabilities were averaged over 100 predictions for each scan and zone. Other top submissions included additional steps, such as pre-processing to normalize the person's height or labeling the dataset's images according to the subject's gender.

4.3.2 ATD in Passive Images

For the ATD model tasked with detection in the passive IR dataset, some of its specific improvements could be the following:

- The IR camera provides a colorbar indicating the temperature distribution in the image, which has been cropped out before inputting into the CNN (as shown in Fig. 4.28). This colorbar has been erased to avoid possible cases in which the CNN has learned to focus on the shapes or numbers of the colorbar, instead of the scene. Nonetheless, the information itself provided by the colorbar can be useful in order to detect when abnormally cold or hot objects are present. Furthermore, this adequately providing this information during learning could allow the CNN to learn how to estimate threat temperatures during testing from new images that do not have the colorbar included.
- In relation to the colorbar, no image pre-processing to normalize temperatures throughout the images in the dataset has been applied. Although the room is stabilized to an air-conditioned temperature, the subtle variations in temperature from threat, clutter, or the different subjects' bodies modify the specific range the colorbar encompasses. A correct way to handle these variations would be to normalize the RGB channels, which correspond to the temperature distribution, in all images before inputting. A related technique would be to convert the image to grayscale and eliminate the RGB channels

altogether, as most public IR datasets seem to do [275,276]. Other authors have also proposed additional pre-processing such as calibration of the IR cameras non-uniform response throughout its pixel array [280], or extraction of object signatures from the images [298].

- As has already been seen in Section 4.2.2, the performance of the model degrades when including unseen variations in the input images during testing. To solve this issue, it is necessary to expand the dataset and take into account new scenarios which are likely to happen in real-life screening checkpoints. Within these, modifications such as more threat classes, multiple threats in a single subject, non-threatening concealed objects, multiple subjects in the image, threat and subject occlusions, or range displacement in the subject, should be included. Depending on the final size of the dataset, traditional data collection and labeling methods might not be appropriate for these expansions, and alternative strategies should be analyzed [237]. In relation to these tasks, making the collected dataset publicly available to the scientific community would provide further insight into better detection approaches and ways to enhance the dataset.
- 3-D surface reconstruction from stereo IR imaging with two or more cameras is possible [277,281]. Such a set-up would allow generating a 3-D image with temperature information from the scene, which in a hybrid imager could be matched to the 3-D range or returned power information provided by an active sensor.

Chapter 5

Other Applications

In this chapter several other applications for the 300 GHz and 122 GHz FMCW radars, aside from security screening, are proposed. Two main areas are addressed with these applications: non-contact evaluation of materials, and vital sign monitoring of human subjects. The chapter includes introductions to these applications, as well as initial evaluations of implementations with both radars. The text is organized as follows:

- Section 5.1 details an application related to non-destructive testing (NDT) for both radars. The use of penetrating mm-wave transmissions allow the radars to inspect various materials in a non-contact manner. This has been applied to a data acquisition technique from air cavities in plastic guides, initially proposed for vehicle guidance, but suitable for any automated data reading requirements. In Annex F, the set-ups and measurement approaches described in this section have also been applied to plastic and fabric material evaluation, analyzing the variations of different materials in the reflected signals to determine whether the radar's are capable of distinguishing among them.
- In Section 5.2 the vital sign monitoring applications is described. The great range resolutions and phase stabilities of both radar systems allow measuring very small range displacements in standoff targets. In vital sign monitoring, this ability is used to measure the chest movements in a human patient due to their respiration and heart beat. This section of the chapter details several of these measurements for both radars.
- Lastly, Section 5.3 presents future improvements and ongoing work for both radars in relation to the two proposed applications. Additionally, two new applications are suggested.

5.1 Data Acquisition on Plastic Guides

A historically common and successful application for microwave technology has been the NDT of objects, meaning the use of material-evaluation techniques without causing damage to the inspected item [57]. NDT techniques usually consist on transmitting an interrogating signal, and analyzing its response after interaction with the tested material. Although NDT can be performed by using alternative technologies, such as acoustic sensing [299] or dye penetrant testing [300], microwave imaging techniques can be specially useful when evaluating objects made from dielectric materials, where electromagnetic radiation can present substantial penetration capabilities [301]. Every method of testing presents different

realms of applicability: for instance, ultrasonic signals may not present good penetration in some materials, and current-inducing techniques may not be useful in lossless dielectrics. More recently, NDT techniques have also been successfully extended to the mm- and submm-wave range [302]. Within NDT imaging methods, those based on FMCW radars have proven to be applicable through multiple frequency bands, and with diverse materials [303–305]. With this in mind, the main advantages of RF-based NDT can be summarized as follows:

- It is based on non-contact measurements, which enables a less intrusive evaluation of fragile or sensitive materials.
- It can use monostatic or one-sided set-ups, based on back-scatter techniques. This reduces the complexity and size of the measurement infrastructure, allowing simpler testing.
- It is relatively inexpensive, and can be performed in real-time with compact and low-power sensors. This has a special importance when taking into account the adaptability of these techniques to field measurements or commercial applications.
- It allows obtaining imaging results, which present an informative and user-friendly alternative to traditional graph- or numerical-based outputs.

In this section an NDT-related application for the 300 GHz imaging radar and the 122 GHz MMIC-based sensor is discussed. The application deals with the use of a mm-wave FMCW radar for data reading from pre-recorded low-cost plastic guides. Through this technique, cm-level plastics are used as storage media for a series of bits (coded as gaps inside the plastic). The radar acts as the reading sensor, using movement to continuously scan the guide and detect the bits.

Although it is true that there is substantial previous work in the field of NDT with micro- and mm-waves, with various researchers specializing in evaluating plastic materials [303, 306–309], these analyses tend to emphasize the evaluation nature of the measurements, usually focused on detecting production imperfections or usage wear. What is proposed in this application is the use of intentional flaws (cavities) in these materials to codify information. In this case, it is not of interest to evaluate the durability or integrity of the material, but rather use it as a vessel of data to compete, in some cases, with alternative media.

Although this technique can be useful for multiple data storage and acquisition applications, it was initially developed for the field of autonomous driving (through vehicle localization) [310]. The development of this method has been motivated by several disadvantages of current commercially-available positioning techniques. For instance, satellite-based methods, such as Global Positioning System (GPS), require a direct line-of-sight to at least three satellites (which may not be possible in indoor locations, overpasses, or tunnels) [70]. Additionally, these systems can be susceptible to jamming techniques, which can cut off communication and positioning [311]. Optical devices, such as high-definition cameras, can also be used to detect lanes, rail markings, traffic signs, pedestrians, or other markers [312], and with a sufficiently accurate pre-composed map of the surroundings, a high-end computer can associate these markers with local positions [313]. Nevertheless, these sensors have limited functionality under common atmospheric conditions such as fog, smoke, or even nighttime [314]. Recent laser-based alternatives, such as Light Detection And Ranging (LIDAR), offer precise local mapping of a vehicle's surroundings [315], although, as was the case with camera techniques, these sensors can be hampered by atmospheric conditions [316]. Additionally, this method also requires complex image processing to generate the desired environment maps, and can be interfered by other nearby LIDARs [317].

The proposed technique, when applied for vehicle localization, means to place a mm-wave FMCW sensor mounted on a vehicle's underbody [71]. As shown in Fig. 5.1, the radar scans

the ground beneath as the vehicles moves, reading the plastic guide installed in the road (for cars) or railway (for trains). An accompanying low-power processor can be programmed to read the bits from the received spectrum. Contrary to LIDAR techniques, where the sensor's processor must differentiate amongst multiple shapes in the scene, in this method the device only needs to determine the range of each gap in the guide (seen as a peak in the spectrum). By recording a binary alphabet, the information in the guides can be associated to whatever data are desired, such as the longitude and latitude coordinates of that spot. For the initial development in this thesis, the objective has been to demonstrate the FMCW sensor's ability to detect cm-level gaps in low-cost plastic guides of different designs, detached from any specific application and focused only on data acquisition.

Figure 5.1: Diagrams of (a) the sensor as it moves through a road with a guide installed, and (b), (c) of the plastic guide with a codification based on two levels of cavity depth. The sensor's movement (vector V) is used to scan the guide. The set-up's axes, used throughout the text, are also depicted.

A small summary of some recent micro- and mm-wave NDT applications, with a special emphasis on imaging systems, can be found in Table 5.1.

Guaranteeing machine-readable data is a crucial aspect of automated industrial systems. This may, for instance, be commonly seen in the retail and food industries where these data are often used to label and identify products within the supply chain [327–331], as well as in other sectors such as wood tracking [332], waste management [333], or parcel handling [334]. For these applications, the two most prominent labeling media and data carriers are barcodes¹ [335] and RFID (Radio Frequency Identification) tags² [336], although less-common alternative devices, such as optical patterns³ [337] or thermal markers⁴ [338] may be used.

While barcodes are still the most-used labels, recent advances in RFID have motivated industrial agents to update their labeling technologies [339]. Although RFID tags initially promised a high return on investment, it is still unclear whether the technology will completely overtake conventional barcodes [340–343]. Cost, tag durability, and technological comprehension are some of the limitations slowing down RFID growth.

A new data acquisition technique based on high-resolution radar technology has been tested on a 300 GHz FMCW radar transmitting 27 GHz of bandwidth, and also on a 122 GHz FMCW MMIC-based radar transmitting 15.1 GHz. As can be seen in Fig. 5.1, a high-frequency radar scans a low-cost plastic guide, with drilled air cavities. The holes can have multiple levels of depth within the plastic, which may be interpreted as bits or symbols. Holes are detected due to reflections from the plastic-air interface created at the boundary of the two media [303].

¹One- or two-dimensional black-and-white patterns that generate an unique electrical signal in the reader.

²Small circuits that transmit an electromagnetic signal that can be picked up by a nearby sensor.

³Visual patterns read by conventional cameras.

⁴Data stored as temperature variations in a label.

Table 5.1: Summary of recent imaging-based micro- and mm-wave NDT applications

Reference	Technique	Material Evaluated	Application	Frequencies (GHz)
Our work (Section 5.1)	FMCW imaging	Plastics	Data acquisition	(286 - 313) & (116 - 131)
Nezadal <i>et al.</i> [318]	SAR* imaging	Plastics	Defect evaluation	(220 - 325)
Nezadal <i>et al.</i> [319]	Polarimetric imaging	Plastics	Defect evaluation	(75 - 110)
Weatherall <i>et al.</i> [320]	Holographic imaging	Explosives	Explosive identification	(10 - 40)
Gao <i>et al.</i> [68]	SAR* imaging	Pigskin	Skin burn diagnosis	(50 - 70)
Laviada <i>et al.</i> [321]	SAR* imaging	Plastics	Pipe evaluation	(8 - 12)
Zoughi <i>et al.</i> [322]	SAR* imaging	Leather	Marker identification	(88 - 91)
Gallion <i>et al.</i> [323]	SAR* imaging	Metals	Crack detection	(26 - 40) & (75 - 110)
Mehdipour <i>et al.</i> [324]	SAR* imaging	Mortars	Steel fiber evaluation	(2 - 4) & (8 - 12)
Friederich <i>et al.</i> [325]	FMCW imaging	Plastics	Radome inspection	(70 - 110) & (110 - 170)
Guillet <i>et al.</i> [326]	FMCW Imaging	Paintings	Painting diagnostic	(75 - 110) & (220 - 330)

*Synthetic Aperture Radar.

Mm-wave radiation possesses electromagnetic properties allowing the previously mentioned application. The large bandwidth available for transmission gives mm-level range resolution (Z -axis in Fig. 5.1), and the small wavelengths permit centimeter-level cross-range resolution (X - and Y -axis in Fig. 5.1) with reasonable antenna sizes. Mm-wave penetration through various materials assures that the system will work through night, fog, smoke, and dust conditions [344, 345], as well as sufficiently probe a plastic guide [346]. Additionally, thanks to the large bandwidth transmitted, complete jamming of the technique is unlikely [24]. These radiation properties hold true for both the 300 GHz and the 122 GHz radars. Although lower, the MMIC-based radar still transmits a large enough bandwidth, with small enough wavelengths, to enable centimeter-level resolutions in all dimensions. Additionally, dielectric penetration capabilities increase as transmitted frequency is lowered, allowing gap-detection further down the plastic.

The proposed technique's main advantage is the durability of the tags, which can be made of diverse materials such as plastics or ceramics, in outdoor environments. Additionally, with the information being inside the material, occasional wear and dents on the guides will not hinder acquisition, and because it is not plainly visible or accesible further privacy is achieved. On the other hand the main disadvantages of the technique are the current cost of mm-wave devices, greater than more mature technologies such as barcodes, and the size of the sensor. However, recent advances in MMICs forecast compact and low-cost sensors [204]. This is where the use of the 122 GHz sensor holds special relevance, because compared with the 300 GHz radar, the new sensor can benefit from lower cost, size, and power consumption, as well as being made completely from commercially-available components (thereby reaching another step towards a future real-world implementation). The volume of the guides can also be problematic: due to centimeter-level spatial resolution, a single bit requires an area of about 1 cm^2 . This could be overcome by packing more information in the Z axis, thereby increasing information density. A simple comparison among the main labeling techniques can be found in Table 5.2. Nonetheless, supposing a 1 cm-long bit (l_{bit}), the proposed technique could still be applied for vehicles moving at highway speeds ($v \leq 120 \text{ km/h}$) and beyond, even when integrating multiple chirps per bit (n) for redundancy, as illustrated in Fig. 5.2 from

$$v = \frac{l_{bit}}{T_{chirp} \cdot n} \quad (5.1)$$

Figure 5.2: Maximum vehicle speed for the proposed technique as a function of bit size (l_{bit}), chirp time (T_{chirp}), and number of chirps integrated per bit (n).

Table 5.2: Comparison of some of the current main labeling techniques

Technology	Durability	Jamming	Privacy	Simultaneous tags [†]	Range	Cost	Line of Sight
Proposed technique	High	Low	High	Yes	Large	High	Required
RFID [336, 340, 341, 347, 348]	Low	High	Low	Yes	Large	Low	Unrequired
Barcode [340, 341, 348]	Low	Low	Low	No	Small	Low	Required
Quick Response (QR) code [348]	Low	Low	Low	Yes*	Small	Low	Required
Near-field Communication (NFC) [349]	Low	High	Low	No	Small	Low	Required

[†]Possibility of reading from multiple tags at once.
*Theoretically possible, but not implemented in available applications [350].

For this application, the 300 GHz radar works on a monostatic configuration illustrated in Fig. 5.3, modified from the design described in Section 2.1.1. In the original set-up, the RF components of the radar are accompanied by multiple metallic reflectors (up to 60 cm in diameter) to focus radiation at 8 m distance. Since for this technique small range distances are the norm, and a reduced size is desired (for possible handheld scanners or vehicle-mounted sensors), a different focusing scheme has been chosen. Two dielectric lenses are placed between the radar's horn antenna and the Object Under Test (OUT). This concentrates the radiated beam onto a spot size small enough to focus on a single bit or symbol in each measurement. A diagram of this scheme is presented in Fig. 5.4. Menlo Systems TPX50 (with 5 cm of focal length) and TPX200 (with 20 cm of focal length) lenses are placed to collimate and focus the radiation, respectively (Annex E).

Figure 5.3: Block diagram and photograph of the RF architecture of the radar used during measurements.

Figure 5.4: Diagram (a) and photograph (b) of the focusing scheme for the proposed set-up. Dielectric lenses are used to focus the transmission onto an area small enough to cover only a single bit. F_1 and F_2 are the lenses' focal lengths (5 cm and 20 cm, respectively).

As for the 122 GHz, the radar is used jointly with the High-Density Polyethylene (HDPE) lens provided with Silicon Radar's kit. The use of this lens eliminates the need for the TPX50 and TPX200 dielectric lenses used for focusing during the 300 GHz testing. Although these lenses have been eliminated, the measurements in this section still use the same working range as the 300 GHz tests (32.5 cm, which takes into account the TPX lenses' focal distances) in order to offer comparable results. In contrast with the dual lens set-up of the 300 GHz radar,

the single lens of the 122 GHz device acts simply as a beam-collimating element (thus limiting the beamwidth at the OUT's range). Because range resolution is critical for the intended application in order to accurately separate bits in the guides, only results from the MMIC configuration with the largest bandwidth (**Configuration 3**, with 15.1 GHz of transmitted bandwidth) are presented⁵. All measurements are performed with a 150 μs chirp. A diagram of the focusing scheme and a photograph of the set-up for the 122 GHz radar are shown in Fig. 5.5, which uses the same coordinate system mentioned in Fig. 5.1.

Figure 5.5: Diagram (a) and photograph (b) of the radiation scheme for the proposed set-up. The MMIC-based radar's HDPE lens is used to collimate radiation on the OUT. The distance to the OUT is set to 32.5 cm, which is the working range taking into account the focal lengths of the lenses in Fig. 5.4, to offer comparable results.

5.1.1 Radar Performance

The signal processing steps used during all measurements, for both systems in this section, are as follows:

- **Signal Calibration.** This stage is performed for the 300 GHz radar only, as it is unnecessary for the MMIC-based radar. The received spectrum must be calibrated to compensate non-linearities from the RF components, as was also done in Section 2.2. To do this, each received chirp is multiplied by a reference signal from an ideal point-like target (a small metallic sphere) at the working range. The effect of performing this calibration on the received spectrum from a flat metallic sheet, and the set-up used, may be seen in Fig. 5.6. As can be seen, the sinc's side lobes have been considerably reduced, which would allow detection of previously masked reflections. The loss of power produced by the calibration procedure can be compensated during processing.
- **Windowing.** To increase the dynamic range and enable the detection of low-power reflections, the received chirps are also windowed with a Hamming function. This is done at the cost of widening the signal's main lobe in the frequency domain, which reduces the effective range resolution (thus forcing an increase in the bits' depth size).
- **Averaging.** In each measurement, multiple received chirps are averaged to form a single signal. The typical measurement and chirp times are ~ 1 s and 150 μs , respectively, producing 6666 complete chirp signals per measurement. Since most measurements are performed on static targets, the chirp signals can be averaged without loss of positioning information.

⁵The two remaining configurations have also been evaluated, but have been found unsuitable in some scenarios.

- **Fourier Transform.** An FFT is performed per chirp to represent the resulting spectrum. In this stage, both the module and phase of the spectrum are of interest.

Figure 5.6: (a) Comparison of uncalibrated, calibrated and ideal spectrums received from a metallic sheet, as captured by the 300 GHz radar. (b) Set-up used for the measurement of the calibration's reference signal from a metal sphere.

An SNR evaluation of the sensor has been performed. The results can be seen in Fig. 5.7 for both the 300 GHz and the 122 GHz radars. In this scenario, the measurement of a commercial mm-wave absorber (as a no-target scenario) has been overlapped with the measurement of a metallic sheet and a solid slab of the chosen guide polymer, HDPE (Annex E), all at 32.5 cm (the lenses' focal distances plus 7.5 cm of variable length between them) in range from the antenna. HDPE has been the chosen polymer due to its properties, such as recyclability, large strength-to-density ratio, and common use as filament for 3-D printing [351], which provide a durable, low-cost and easy to machine material.

Figure 5.7: Measurements at (a) 300 GHz and (b) 122 GHz (**Configuration 3**) of a metallic sheet, a solid slab of HDPE, and a mm-wave absorber. RX bandwidth is 20 MHz, sweep period is 150 μ s, and a Hamming window is applied.

Considering the metallic sheet as an ideal target, the maximum attainable SNR and the

maximum noise peak have been measured in both figures, where range gating is applied around the working distance (32.5 cm). For the 300 GHz radar, the returned power for HDPE is 12 dB lower than for the metal sheet, due to transmissivity through the plastic and surface roughness which scatters the signal, but high enough to be separated from the maximum noise. A prominent first peak can be detected at 33 cm and a second one, just as powerful, at 41.2 cm. These are due to the front and back surfaces of the slab, respectively, confirming that the sensor is capable of penetrating the HDPE. As for the 122 GHz radar, the HDPE also reflects lower power than the metallic plate, although the peaks from both of the plastic's surfaces can still be resolved from the noise (confirming the radar's penetration capabilities). The multipath peaks present in the 122 GHz measurement only when illuminating the metallic sheet are due to reflections between the sheet and the radar's PCB. This can be concluded from the progressive doubling of their range (i.e. the sheet appears at 0.3 m, the first multipath reflection appear at 0.6 m, and so forth). Although not visible, the HDPE slab can also create multipath clutter, but its less-powerful reflections fall below the noise floor. A simple solution to this effect is the reduction in size of the radar's PCB, which acts a reflector to the back-propagation from the target.

The illuminated XY spot size at the focused distance has been measured, to adapt guide and bit dimensions to the appropriate sizes. To determine this parameter, the criterion has been the return loss. Assuming that the OUT has similar roughness, material, and orientation as the metallic sheet, a loss of reflected power indicates that the sensor's illumination area is larger than the target. Coins have been used as an area-varying target due to their similarity to the reference metallic sheet. The results can be seen in Fig. 5.8. For the 300 GHz system the 1€ coin (diameter $\varnothing = 23.25$ mm) reflects as much power as the sheet, whereas with the 0.1€ coin (diameter $\varnothing = 19.75$ mm) the returned power is 1 dB lower, indicating that the XY spot size is between their two diameters. Meanwhile, the resolution in the Z axis is 1 cm, given by (1.20), when $BW = 27$ GHz and $\alpha = 1.81$ from Hamming windowing. For the 122 GHz radar, the results present just 0.5 dB losses when illuminating the 1€ coin, similar to the measurements at 300 GHz, but 14 dB when targeting the 0.1€ coin. This would indicate that the spot size for this MMIC is between their two diameters, but much close to $\varnothing = 23.25$ mm than $\varnothing = 19.75$ mm, in contrast to the 300 GHz system.

Figure 5.8: XY spot size measurement at (a) 300 GHz and (b) 122 GHz (**Configuration 3**), with a return loss criteria, and using metallic coins of various sizes as targets.

Due to body movement, in applications based on handheld-scanners a constant distance to the guide cannot be assumed. Similarly, when mounted on a vehicle, the suspension system can sometimes cause variation in the distance from the underbody to the ground.

Since a stable length from the sensor to the guide may not always be assured in a real-life scenario, the system requires a range tolerance that includes possible variations. To measure the system's tolerance, a 1€ coin has been used. Starting from the focused distance (32.5 cm) the coin is moved 10 cm back and forth, in 1 cm steps. The recorded returned power is shown in Fig. 5.9, compared against noise floors with the maximum levels seen in Fig. 5.7 for both radars.

Figure 5.9: Range measurement sequence at (a) 300 GHz and (b) 122 GHz (**Configuration 3**) with a 1€ coin moving 10 cm back and front starting from 32.5 cm, in 1 cm steps, to analyze tolerance to variations of target placement in the Z axis.

At 300 GHz, as can be seen in Fig. 5.9(a), while the Z distance is modified the maximum received power decreases slightly, but still stays more than 24 dB above noise in all cases. This stability in returned power is provided by the tolerance of the focal scheme. Meanwhile, at 122 GHz in Fig. 5.9(b), as the range is increased, the reflections become less powerful and near the noise floor due to spillover from the illuminating spot becoming bigger than the target as the lens's collimation degrades with distance.

Another important aspect in the radars' operation is their orientation in the Z axis. Due to the small wavelength of the frequency band, specular reflection can cause losses in returned power when there is non-normal incidence. The losses for various orientation shifts have been measured, for the HDPE slab at 32.5 cm, and shown in Fig. 5.10.

For the 300 GHz system, shifts up to $\pm 2^\circ$ cause little effect, but shifts of $\pm 5^\circ$ or more lower the returned power to noise floor levels, rendering the sensor unable to properly detect the HDPE. A solution could be using rounded shapes (instead of flat surfaces), which provide better independence from incidence, for bits and guides. Comparing these results to those from the multistatic imaging section [Fig. 2.38(c)] it can be seen that the losses in this new measurement are much smaller for the same rotations. This is due to the effect of the target's standoff range [just 32 cm in the new measurement, but 8 m in Fig. 2.38(c)]. Comparing the 300 GHz and 122 GHz results, it can be seen that the 122 GHz radar suffers from less loss when encountering the same shift. For instance, at a $+1^\circ$ rotation the MMIC-based sensor is affected by, in the worst scenario, a 0.21 dB loss, while for the 300 GHz it is a 0.71 dB attenuation. This could be due to the increase in transmitted wavelength (where the effect of specular reflection is less noticeable for the same target) and to a narrower generated beamwidth.

Although THz radiation can propagate in various atmospheric conditions, water presents high absorption in this band. Thus, fog, rain, and puddles over the guide can cause vanishing

Figure 5.10: Measurement of losses due to the sensor's orientation shifts in the Z axis at (a) 300 GHz and (b) 122 GHz (**Configuration 3**). A 0° shift means the target's surface normal aims directly at the radar. For the 122 GHz radar, the measurements performed include data with and without its dielectric lens.

of the information. To analyze the sensor's susceptibility to different humidity scenarios, the spectrum of a solid HDPE slab has been measured under various surface-wetness conditions. The results are shown in Fig. 5.11.

Figure 5.11: Spectrum of solid plastic slab in different humidity scenarios at (a) 300 GHz and (b) 122 GHz (**Configuration 3**). In the *wet* scenarios water is applied with a diffuser on the surface. In the *film* scenario plastic film is placed over the wet surface.

In the 300 GHz case it can be seen that applying water with a diffuser (simulating a light drizzle) does not significantly affect detection. This can be due to two reasons: the water tends to slowly slide off the surface, and the use of a diffuser does not ensure uniformity of the water applied (i.e. the focused spot might not be fully wet). On the other hand, when the water is trapped with wrapping film (simulating a small pond over the guide), the liquid becomes an absorber. This is specially noticeable for the back peak in the *film* scenario (where the applied water is trapped with film over the HDPE's front surface), which is undetectable from the noise, since penetration through the plastic becomes harder due

to the water's absorption at the surface. To overcome this obstacle and displace water in a future commercial implementation, the guide can be coated with hydrophobic paint, designed with a curved surface, or the sensor accompanied by a small fan. At 122 GHz, in Fig. 5.11(b), the received power from the plastic in both wet scenarios decreases when compared against the dry slab. Nonetheless, both peaks can still be resolved from the noise, in contrast to the results at 300 GHz. This is because, just as material penetration improves as the transmitted frequency is lowered, the absorption by water molecules is decreased⁶ [16]. These results seem to indicate that the 122 GHz sensor would be better suited for this application than the 300 GHz radar, under water conditions.

Because a potential application for the technique is vehicle localization in conventional roads, the effect of asphalt reflections in the sensors has been tested. These reflections can be present in areas where no guide is installed, or when the sensor is misaligned. The result of measuring the spectrum received from the asphalt at the working range previously used (32.5 cm) is presented in Fig. 5.12.

Figure 5.12: Spectrum of a block of road asphalt at (a) 300 GHz and (b) 122 GHz (**Configuration 3**), compared against the received reflections from a metallic sheet, a solid HDPE slab, and a mm-wave absorber.

At both 300 GHz and 122 GHz, the power levels registered for the asphalt scenario are similar to those from the mm-wave absorber, and much below the peaks of the metallic sheet and plastic⁷, indicating that almost no reflections are received by the sensor. This is likely due to the rough surface of the asphalt block, which scatters the transmitted signal in multiple directions [352]. These results would indicate that the sensors are capable of accurately distinguishing the different materials.

The performance characterization done with these experiments would indicate that both radars are similarly apt for the data acquisition application proposed. Both sensors presented comparable SNRs for the HDPE material and performance in the asphalt scenario. Specifically, each sensor has its own advantages. While the 300 GHz system presents better resolutions in range and cross-range, as well as better tolerance to displacements in the range dimension due to a focusing scheme, the 122 GHz sensor performs slightly better under rotation and water scenarios. Of course, the reduced cost, size, and power consumption of the 122 GHz device must also be taken into account.

⁶ According to Fig. 1.2, the absorption coefficient would be $\sim 50\%$ greater at 300 GHz than at 122 GHz.

⁷ At 122 GHz the asphalt falls 20 dB below the plastic, while at 300 GHz the asphalt is 35 dB lower than the HDPE.

5.1.2 Guide Measurements

To validate the proposed technique in regard to its bit reading capabilities, a plastic guide with rectangular gaps has been designed, seen in Fig. 5.13(a). The first row of air gaps is at a fixed Z distance, while the second row is progressively further away from the first gap. Information can be coded, even with bases greater than 2, using the distance from the sensor to the second row, the distance between gaps in a column, or even just the existence of a gap in a given row. Information density, limited by sensor resolution, can be incremented with the use of 2-D and 3-D codes and appropriate scanning patterns (although in this initial stage only 1-D code designs have been considered). CAD diagrams for the guides designed can be found in Annex E.

Figure 5.13: Photograph of the first two 15x2.5x5 cm³ HDPE guides designed, with (a) internal-gap and (b) surface information.

The spectrum of the bit columns in Fig. 5.13(a) is seen in Fig. 5.14 at 300 GHz, and in Fig. 5.15 at 122 GHz (in this case only for shown for the first column⁸). As can be seen in Fig. 5.14, the measured range between the front and back surfaces is near 7 cm, which is 2 cm more than the real distance in the guide. This is due to the variation of electromagnetic propagation velocity as a function of the medium (in this case, air and plastic). Knowing the total depth of gaps (L), guide thickness (T), and the propagation speed in HDPE ($c = 0.66 \cdot c_0$) [353] the distance estimate of any reflection received by the radar may be corrected by:

$$R_r = \left(\frac{R_u}{T}\right) \cdot (L + 0.66 \cdot (T - L)) \quad (5.2)$$

where R_u is the uncorrected estimate, and R_r the real distance. Substituting 5 cm into T , 7 cm into R_u , and 1 cm (two 0.5 cm gaps) into L , an R_r of 5.09 cm is obtained, which coincides with the guide thickness. The distance errors in Fig. 5.14 have been calculated after correcting this problem, and all columns present similar error margins. As expected, while the second row of bits grows closer to the back surface peak separation becomes more difficult⁹.

In Fig. 5.15 the 122 GHz sensor is also able to separate all four peaks¹⁰, although with less resolution than the 300 GHz case. As was the case in Fig. 5.14, the range measurements must be corrected by (5.2) due to the signal's propagation through the HDPE.

⁸The remaining three show similar results.

⁹Specially noticeable for the last column, where the distance between the back surface and the second gap is minimum, in Fig. 5.14(d).

¹⁰This does not happen if **Configuration 1** or **Configuration 2**, which use less transmitted bandwidth, are used.

Figure 5.14: Spectrum of the bit columns in Fig. 5.13(a) at 300 GHz. Columns are numbered according to their position in Fig. 5.13(a), starting at the left side of the photograph. The numbers in parenthesis are the error, produced by Z-axis resolution and FFT precision, in the measurement of distance for each surface.

Figure 5.15: Spectrum of the first column (starting from the left) in Fig. 5.13(a), for **Configuration 3** of the 122 GHz radar.

Continuous movement sequences, with complete guide scanning (as in a real-world data-reading application) have also been performed. In these sequences, the plastic is measured every 5 mm in the X axis. The result of applying this sequence to the solid HDPE slab with no cavities may be seen in Fig. 5.16. For the 300 GHz radar, as expected, both surfaces are detected throughout the entire sequence with little variation in range or power. The back surface presents less powerful reflections due to propagation losses through the plastic. Since no cavities are placed in the slab, no reflections are detected inside. For the 122 GHz sensor, it can be seen that it is also capable of separating the two surfaces of the slab and the results are equal to the 300 GHz radar (although with wider surfaces due to the lower range resolution).

Figure 5.16: Continuous scan of a solid HDPE slab at (a) 300 GHz and (b) 122 GHz (**Configuration 3**), with a -50 dBm threshold. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.

The results of the same measurement sequence, applied to the guide design shown in Fig. 5.13(a) may be seen in Fig. 5.17.

Figure 5.17: Continuous scan of the guide in Fig. 5.13(a) at (a) 300 GHz and (b) 122 GHz (**Configuration 3**), starting at the leftmost column, with a -50 dBm threshold. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.

At 300 GHz, the front surface of the guide (at 33 cm) is homegenously detected, while

the back surface (at 41 cm) has more variation in power and range. For the power variations, this is because the penetration and reflection losses of the HDPE result in a less powerful reflected signal at the back surface. On the other hand, the range variation (seen in X positions where there are gaps) is produced by the different media through which the signal travels, as indicated by (5.2): when there are gaps, the waves travel less distance through HDPE, and thus the range estimation from vacuum velocity holds truer. The first row of gaps (at 36 cm) and the second row (starting at 38 cm), as well as the second's growing separation, are also detectable. For the 122 GHz radar, the sensor is only capable of properly separating only for the first row of gaps, which presents larger separation to the guide's surface than the second. This guide design could still be used for the 122 GHz device, but information could then only be coded into that first row.

A second guide, using surface information instead of gaps, is shown in Fig. 5.13(b). In this case, the information is in the front-to-back surface distance, instead of inside the guide, where penetration losses occur. With this design information is less dependent of received power and bit dimensions can be larger. The results of the continuous measurement sequence for the second guide may be seen in Fig. 5.18. In the 300 GHz case, the back surface is detected with more power than the front due to the elimination of the multiple reflections caused by the air gaps, and bits with the same value and in adjacent position are correctly detected as being twice as long as an individual bit. For the 122 GHz radar the results improve, compared to Fig. 5.17(b), due to the increase in bit dimensions¹¹.

Figure 5.18: Continuous measurement sequence for the guide in Fig. 5.13(b) at (a) 300 GHz and (b) 122 GHz (**Configuration 3**), starting at the rightmost column and with the bits at the back surface, with a -50 dBm threshold. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.

An alternative third guide has also been designed, as shown in Fig. 5.19(a). In this case, the 4 bits are codified by triangular cavities at varying depths within the guide. The triangular shape has been chosen for its similarity to corner reflectors used in conventional radar systems as calibration targets, offering powerful reflections. The guide has been scanned in continuous movement in two different orientations as illustrated in Fig. 5.19(b), depending on the surface that the radar faces. For the first orientation, the radar faces the top vertices of the triangular bits, while in the second orientation the radar faces the triangles' bases. The results for scanning in both orientations at 300 GHz may be seen in Fig. 5.20, and in Figs. 5.21 for the 122 GHz radar.

¹¹For this guide design the reduced-bandwidth configurations of the MMIC-sensor (**Configuration 1** and **Configuration 2**) are also viable alternatives.

Figure 5.19: Photograph of an HDPE guide based on triangular gaps as bits and diagram of the radar orientations considered when measuring the guide.

Figure 5.20: Continuous measurement sequence for the guide in Fig. 5.19 at 300 GHz, with a -50 dBm threshold. In (a) the radar is placed in the first orientation, according to Fig. 5.19(b), while in (b) the radar is in the second orientation of that figure. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.

Figure 5.21: Guide in Fig. 5.19 for the (a) first and (b) second orientations considered, for **Configuration 3** of the 122 GHz radar. Range variations due to propagation through different media are uncompensated. Colorbar: Received power in dBm.

At 300 GHz, it can be noted that in the first orientation only the front surface of the guide is detected (without gaps or back surface seen). This is due to specular reflection from the top corner of the bits, which directs the reflected wave away from the RX. In the second orientation, because the travelling wave first encounter the flat bottom side of the bits and thus peaks are detected in the three of the four columns (the fourth column is missing because the triangle has no bottom side). In this last case, the back surface remains undetected because after the bottom side of the triangle the transmission reaches the top corners once again. For the 122 GHz radar, it can be seen that the specular reflection from the top corner of the triangles still produces the vanishing of bits and back surface. Additionally, bits are now harder to separate from each other and the front surface than with the 300 GHz radar. In this case bit separations would need to be larger for this guide design if it is to be used with the 122 GHz sensor.

In conclusion, a new technique for data acquisition has been proposed, using a 300 GHz radar to read data recorded in an HDPE guide. Additionally, the technique initially proposed at 300 GHz has been transferred to a 122 GHz sensor. The 300 GHz radar's ability to detect the gaps in which the information is encoded has been proven. Meanwhile, the implementation at 122 GHz has been completely successful for gap detection only in the case of the MMIC source with larger bandwidth, which indicates that a considerable frequency excursion is needed for the application, but still demonstrates that this parameter (and the central frequency) can be lower than initially designed.

Further work remains in the area of sensor integration via low-cost MMICs (to reduce sensor cost and size), which is where the successful use of the 122 GHz proves to be a benefit for a future industrial application, as commercial components in lower frequency bands and with less bandwidth tend to be available sooner and with lower cost. The promising results of this technique with a MMIC-based radar denote a step forward into more compact sensors; an advantage for both handheld or vehicle mounted scanners. For any sensor chosen, appropriately designing the antenna or focusing element is a priority to enable small spot sizes and thus small bit dimensions. Additional research must also be done in the area of guide design, in order to optimize the size, shape, and distribution of bits. Additional plastic and non-plastic materials can also be analyzed, and potentially used for different scenarios.

While THz technology's current high cost and plastic guide size inhibit item-level tagging,

the proposed technique can be useful in a hybrid system, where it would coexist with other labeling media. Such a system would utilize each labeling media in their optimal placement. Barcodes could be used in item-level tagging, while RFID tags are placed in packages. The proposed application can be more suitably used for container- and pallet-level tagging, where outdoor conditions and wear-and-tear have insignificant effect on the guides, and tagged-item size permit larger guide and bit dimensions. With this implementation (illustrated in Fig. 5.22), the proposed technique could help reduce the cost of lost containers in a supply chain [354]; an obstacle that cost the automotive industry \$1.4 billion in 2005 [355].

Figure 5.22: Hybrid item-labeling system in container storage, where different tagging technologies are used at different storage levels to optimize each technique’s advantages.

5.2 Vital Signs Monitoring

In many medical scenarios, the detection and monitoring of a patient’s vital signs, such as respiratory rate, heart rate, or blood pressure, is a first-order priority. For this purpose, several technologies can be used. These solutions can be divided mainly into three categories: on-body equipment [356], ingestible sensors [357], and non-contact monitoring [358]. On-body equipment is the most common form of monitoring, and can range from simple electrode patches used to track heart rate, to clinical-grade wearable devices that measure oxygen saturation, pulse wave, or quality of sleep. On the other hand, ingestible sensors are a recent innovation, which enable the detection and tracking of a subject’s vital signs and additional features from inside the body, relaying the information to outside RXs. Lastly, non-contact monitoring encompasses various technological solutions based on remotely tracking a patient’s signs, without acting on the subject’s body. Among these three categories, non-contact monitoring poses several considerable advantages:

- **Eliminates the need for skin contact.** Through remote monitoring, the need for placing support devices, such as patches or straps, is eliminated. This is very beneficial in cases where the patient suffers a skin condition that hinders some types of contact, such as burn victims or newborns. In addition, this reduces cost and hygienic concerns, since a single remote sensor can be used in as many patients as needed, without having to discard it or conduct sanitary hazards.
- **No wearable devices required.** The disposal of wearable sensors allows for less invasive and obtrusive measurements, where the patient can perform daily activities and move freely without the discomfort of a constantly-worn item. This is specially beneficial in long-term measurements, such as over-night baby monitoring, where the unpleasantness of wearing the sensor can pile up. Additionally, measurement precision

can be increased, when compared to contact-based solutions, due to the patient's unawareness of the measurement¹². With these benefits in mind, remote-only monitoring can then be performed in locations aside from medical facilities, enabling early detection of medical conditions and in-home monitoring, which can then provide more frequent measurements without imposing higher patient throughput in saturated hospitals and clinics.

- **Enables multi-patient monitoring.** Some non-contact sensors can allow monitoring several patients at once, which obviously provides considerable cost and measurement time reductions. This is, for instance, the case of video-based monitoring systems, which can track data from multiple patients in a room with a single camera.
- **Capable of simultaneously measuring multiple parameters.** Depending on their mode of operation, some non-contact sensors can also measure various medical parameters at once. For example, RF-based systems can be used to simultaneously measure a patient's respiration and heart rate, among other signs.

Within non-contact monitoring, several technologies can be used to measure a patient's vital signs. For example, video-based systems have proven to be able to monitor various parameters, such as respiratory rate, oxygen saturation, or cardiac activity [359]. On the other hand, IR imaging systems are particularly useful for measurements dealing with a patient's body temperature, due to the body's own blackbody radiation [360]. However, sensors based on RF transmissions are perhaps the largest category, and have demonstrated their monitoring capabilities since the 1970s [65], both in animal and human patients, and for multiple applications (such as studies of sleep apnea [361] or sudden infant death syndrome [362]). Although the RF systems used were initially bulky and expensive, advances in MMIC technology have enabled the appearance of integrated sensors in low-cost and small chips. Additional advantages to non-contact RF-based monitoring would be its independence from sensor orientation, its permission for the subject's use of clothing during measurements, and the elimination of the need for sedation to steady patients in a measurement. Furthermore, because the transmission used usually presents low-power radiation levels, safety is not an issue. A summary of recent examples of RF-based remote vital sign monitoring applications can be found in Table 5.3.

Furthermore, RF remote vital sign monitoring sensors are not exclusively present in academia. Various companies offer commercial products for this task, such as:

- Novelda AS, through their XeThru series, presents various radar chips and evaluations boards for respiration detection [374].
- SpecOps Group Inc. offers the 3.15 GHz X3FINDER device, which can detect multiple heartbeats and respirations through materials such as concrete or wood (which can be specifically useful in collapsed buildings, tunnels, and mines) [375].
- Neteera Technologies Ltd. also offers a 140 GHz chip for simultaneous respiration and heartbeat measurement [376].
- Texas Instruments Inc. has also demonstrated their 77 GHz radar chip's ability to monitor a driver's respiration and heartbeat [377].
- Infineon Technologies have used a 60 GHz chip to measure respiration and heart rate, among other variables [378].

¹²Patients aware of being subject to measurement can sometimes involuntarily cause forced respiration patterns.

Table 5.3: Summary of recent RF-based remote vital sign monitoring applications

Reference	Technique	Subject	Application	Frequencies (GHz)
Our work	FMCW	Adult human	Respiratory rate detection	(116 - 131)
Our work	FMCW	Adult human	Heart rate and waveform detection	(116 - 131)
Our work	FMCW	Adult human	Respiratory and heart rates detection	(286 - 313)
Kuo <i>et al.</i> [363]	Doppler + clutter canceler	Adult human	Simultaneous heartbeat and respiratory rate detection	60
Tu <i>et al.</i> [364]	Doppler	Adult human	Respiratory rate detection under 1-D body movement	5.8
Radzicki <i>et al.</i> [365]	Doppler	Adult human	Through-the-wall heartbeat detection	3.9
Kagawa <i>et al.</i> [366]	Doppler (two radars)	Adult human	Sleep stage classification	24
Forouzanfar <i>et al.</i> [367]	PMCW*	Adult human	Respiratory event classification	24
Wang <i>et al.</i> [368]	FMCW+Interferometry	Adult human	Simultaneous heartbeat and respiratory rate detection + 2-D positioning	(5.72 - 5.88)
Huan <i>et al.</i> [369]	Doppler	Rat	Simultaneous heartbeat and respiratory rate detection	60
Singh <i>et al.</i> [370]	Doppler (two radars)	Lizard	Physical activity detection	10.52
Hafner <i>et al.</i> [371]	Doppler (two radars)	Fish	Underwater heartbeat detection	2.4
Huang <i>et al.</i> [372]	FMCW	Infant human	Simultaneous heartbeat and respiratory rate detection	(3.1 - 5.3)
Suzuki <i>et al.</i> [373]	Doppler (two radars)	Elderly human	Simultaneous heartbeat and respiratory rate detection	24

* Phase-modulated continuous-wave.

In addition to commercial products, multiple organizations have hosted international competitions to develop new breakthroughs in vital sign monitoring, displaying the widespread interest in this field. This is the case of Qualcomm's \$10 million Tricorder XPRIZE [379], Children's National \$250k Pediatric Medical Device Competition [380], or Innovation4Health's \$20k Challenges [381].

Among the RF-based solutions, two main sensor categories can be established: Doppler radar [66], and wide-bandwidth systems [382]. In Doppler radars, a continuous wave (CW) of a single frequency is transmitted, and the vital sign information lies exclusively on the phase history of the returned signal. Because no bandwidth is transmitted, these radars have no range resolution, and thus the measurements can suffer from poor range-based isolation of the target from clutter. Additionally, dc offsets and channel imbalances in these systems can cause substantial difficulties, because of their influence on the received signal's phase. On the other hand, wide-bandwidth systems (such as in an FMCW configuration) can provide excellent resolution to separate one or multiple targets, and enhance phase-based information with range data. However, these systems are usually restricted to low-power transmissions, causing the target's SNR to be reduced, and thus possibly affect measurement precision. In both CW and FMCW systems, a high operating frequency is desired to offer increased phase precision from shorter wavelengths in the former, and range resolution from wider bandwidth available in the later.

In the case of the two radar systems presented in this thesis, remote vital sign monitoring is also possible. For the 122 GHz sensor, a successful implementation of this application would demonstrate an initial hand-held, low-cost, and low-consumption prototype. For the 300 GHz system, favorable results would enable a hybrid system, capable of simultaneously obtaining a 3-D image of the patient¹³ complemented with vital sign information. Both systems have been tested in different remote vital sign monitoring scenarios, with the results seen in the following sections.

5.2.1 For the 122 GHz radar

One of the main benefits of the great phase stability of the developed 122 GHz radar is the enablement of measuring small target displacements within a range bin, as demonstrated in Section 3.1. Because for the considered bandwidths range resolution is limited to cm-level values, the use of frequency-only evaluation for vital signs monitoring is not possible. This is because typical amplitudes for chest and subcutaneous tissue movements due to regular respiration and heartbeat lie between 4 - 12 mm and around 0.5 mm, respectively [383], which are smaller than the range resolution. This means that phase evaluation is necessary to accurately detect and track both movement types; the radar must detect the vital signs through the phase modulation produced on the returned signal. Due to the μm -level phase stability of the 122 GHz radar, respiration and heartbeat measurements are possible with the sensor.

To study the 122 GHz radar's potential use as a vital signs monitoring tool¹⁴, several experiments have been performed. First, a shirtless sitting subject is measured with the radar aiming at his chest area, approximately 1 m away. The set-up for this experiment may be seen in Fig. 5.23. The subject is static but breathing normally, and his shirt has been removed to reduce signal attenuation and target movement. The radar transmits 15.1 GHz of bandwidth with 150 μs chirps continuously for 10 seconds. During signal processing, an FFT is applied to each chirp to compute its spectrum. The spectrum is then range-windowed around the subject's known distance, and the first peak detected (corresponding to the subject's skin)

¹³Useful, for instance, to scan for burns, tumors, and other skin-level anomalies.

¹⁴At the time of writing this text, a recent early access publication has emerged, demonstrating some vital sign scenarios for the same 122 GHz MMIC [384].

is selected. As mentioned previously, for vital sign detection, frequency-only measurements with the radar are not sufficient, and thus, for this scenario phase-based tracking of the selected peak's range is added to the processing pipeline, through the relationship in (3.6). The complete signal processing chain for the measurements may be seen in Fig. 5.24. With this processing, the detected peak's associated range is calculated and tracked throughout the 10 second measurement.

Figure 5.23: (a) Photograph and (b) diagram of the measurement set-up for vital sign detection and tracking with the 122 GHz radar.

Figure 5.24: Block diagram displaying the signal processing stages for vital sign monitoring with the 122 GHz radar.

The results of applying this processing chain and set-up to two scenarios with different breathing rate can be seen in Fig. 5.25, where the radar has been capable of accurately following the chest's movements. The sinusoidal shape of the resulting data is due to the thoracic cavity's inflation and deflation during a regular breathing cycle, which causes the chest's distance to the radar to vary accordingly. Because each breathing cycle is independent, so is the increase it creates in the thoracic cavity's volume, and so the measured data's amplitude is not constant. The radar has no problem in accurately representing the two different breathing rates, because their frequency (approximately 0.25 Hz and 0.6 Hz, respectively) is much lower than the maximum unambiguous frequency detectable. In these measurements, the calculation of the maximum unambiguous frequency is analogous to the one derived from applying the Nyquist criterion to the sampling frequency in a conventional ADC system. In our implementation, sampling is dictated by the chirp period, since each chirp corresponds to a single distance measurement. For the 150 μ s chirp used during the measurement, the equivalent sampling frequency would be 6666 Hz, and thus the maximum unambiguous frequency would be 3333 Hz (which is well above the detected respiration rates).

Figure 5.25: (a) Measured chest displacement during the subject's low breathing rate scenario. (c) Measured displacement during the subject's high breathing rate scenario.

The radar can also detect the specific rate at which the breathing is taking place. This is done by adding the final FFT stage of the signal processing chain. The Fourier transform is applied to the distance fluctuations from a measurement and a peak-detection algorithm is run on the spectrum. Because captured respiration signals are sinusoidal, the spectrum will feature a prominent peak at the signal's frequency, which corresponds to the subject's breathing rate. An example of this application, for a measurement scenario equal to Fig. 5.23(a) is presented in Fig. 5.26. The spectral peak detected registers a 0.8 Hz breathing rate, which is in perfect agreement with the time data (where 8 respiration cycles are detected in 10 seconds), and at the upper limit of regular respiration rates (usually between 0.1 - 0.7 Hz [385]).

Figure 5.26: (a) Measured chest displacement during the subject's breathing. (b) Spectral representation of the measured data.

Having confirmed the radar's ability to detect respiration signals, the next experiment deals with its ability to detect the subject's heart rate. For this experiment, the measurement set-up in Fig. 5.23(a) is repeated, although the subject is now asked to hold his breath during the measurement. This is done to avoid masking the small displacements due to heartbeat

with the stronger movements of the respiration cycle. Because the heartbeat is not always visible on the subject's surface, another form of confirmation must be used to validate the measured data. For this task, a small evaluation board from SparkFun is used. The board includes a heart rate monitoring chip¹⁵, which allows the measurement of the heart's rhythm through an electrocardiogram (ECG). The ECG measures the heart's electrical activity [386], which is associated to the cardiac muscle's polarization, with the use of electrode patches attached to the subject's skin. In addition to heart rate measurement, the ECG enables the representation of finer movements in the cardiac cycle, allowing the detection of various anomalies, such as arrhythmias, seizures, or heart attacks.

As for the respiration scenarios, the signal processing pipeline for the heartbeat detection and tracking is based on combined range and phase evaluation, where the phase of a selected range peak is registered for every chirp signal during the measurement. The resulting data from the radar and the ECG board for this set-up may be seen in Fig. 5.27(a). Notice that the y axis is labeled as simply *Amplitude*, since the ECG signal is measured in volts, while the radar signal corresponds to a distance displacement. As seen in the results, the radar's data synchronizes perfectly in time with the ECG. Each of the ECG's larger spikes (named QRS complex [387]) takes place at the instant where the radar's periodic signal is at its lowest point. This means that the QRS complex is associated with the instants where the subject's skin is closer to the radar, which is the case when the blood vessels are expanding due to heartbeat. This operation is as expected, since the QRS complex corresponds to the depolarization of the heart's ventricles as they contract [388], exactly the moment where the heart pumps blood out through its vessels and generates the subject's pulse (which causes the vessels to expand and grow closer to the radar). The synchronization between ECG and radar data means that the radar is capable of detecting and tracking the subject's heart rate¹⁶.

Figure 5.27: (a) Measured results for a resting subject's heart rate, compared against the ECG board's data. (b) Spectral representation of the radar-measured data.

As for the heartbeat signal, the heart rate can be extracted from the spectrum of the radar's measured data. This is shown in Fig. 5.27(b), where an FFT has been applied to the radar data from Fig. 5.27(a) to detect the spectral peak of the semi-sinusoidal heartbeat.

¹⁵Model AD8232 from Analog Devices.

¹⁶Both the radar and the ECG need to be properly located for the measurements to be accurate. Traditionally, ECG measurements require 6 electrodes to be placed between the fourth and fifth intercostal spaces in the ribcage, and deviation from those locations will produce errors [389]. To be synchronized with the ECG, the radar should also be aimed at those locations.

The main peak detected has a frequency of 0.94 Hz, which would correspond to a heart rate of 57.6 bpm (within the typically-encountered heart rate values of 0.9 - 3 Hz [385]). Meanwhile, a second less-powerful peak is detected at 1.89 Hz, which would correspond to a heart rate of 113.4 bpm. Since this value is above the upper limit of a resting person's rate (usually 1.67 Hz or 100.2 bpm [390]), it cannot be due to the subject's heartbeat, as is most likely a x2 harmonic of the pulse signal (also very common in other forms of heartbeat measurement [391]).

In addition to heart rate information, more data can be extracted from the individual cycles detected by the radar. By isolating a single cycle from Fig. 5.27(a) (between 4 s and 5 s), the radar's data highly resemble the shape of a photoplethysmogram (PPG) [392], which is an optical measurement of the pulse waveform in skin. This comparison may be seen in Fig. 5.28. The PPG allows measurement of the duration and intensity of the systolic and diastolic phases of the cardiac cycle, as well as the detection of the dicotic notch (a change in cardiovascular pressure from the closure of the aortic valve). Thus the data from a PPG provide additional information regarding the correct operation of the subject's cardiovascular system. The resemblance of the radar's data to the PPG shape (in this experiment and other radar measurements in the literature [393]) seemingly indicates that the sensor can also be used for those purposes.

Figure 5.28: (a) Diagram of the pulse's waveform as measured with a PPG, reproduced from [394]. (b) Isolated heart cycle measured with the radar in Fig. 5.27(a).

5.2.2 For the 300 GHz radar

In the case of the 300 GHz system, the tests also include respiration and heartbeat detection, in separated experiments. Although the radar's phase stability is also in the order of μm [127], the large standoff distance for this system increases the measurement uncertainty of the set-up. To limit errors during heart rate measurements (which present the smallest displacements) in this system, a more controlled set-up is used, very different from the usual security screening scenarios commonly studied with the radar. The subject will be holding his breath and sitting down in a chair placed in the FoV, and will also be lightly tied to its back with a seat belt to avoid involuntary body movements. Both clothed and nude measurements have resulted in successful detections. A sample heartbeat measurement from a sweater-clothed young male is shown in Fig. 5.29, when the radar aims at his fifth intercostal space (where the heart's apex is located) for maximum displacement.

Figure 5.29: (a) Measured heartbeat displacement with the 300 GHz radar. (b) Spectral representation of the measured data.

As can be seen for this radar the measured displacement is centered around its standoff range. In Fig. 5.29(b) the peak frequency is 1.202 Hz, which corresponds to a 72.12 bpm heart rate (viable result for a healthy resting male). When compared to the 122 GHz results in Fig. 5.28, it is clear that the 300 GHz system does not reconstruct the pulse waveform as accurately. Since the 300 GHz radar's phase stability is in the same order as the 122 GHz system, it is expected that with a more-controlled set-up the waveform can also be replicated in future measurements. When comparing Fig. 5.29(a) and Fig. 5.27(a) it can also be noticed that they share an anomaly in the 0 - 1 s range. This is an involuntary movement as the subject suddenly holds his respiration at the measurement start.

As for respiration, the breathing signal of a male subject standing at a distance of 8 m has been detected, with the results seen in Fig. 5.30. For respiration detection, because the chest displacements are larger, small involuntary body movements can be tolerated and thus the subject can be untied and standing. The FFT-detected breathing rate is 0.3 Hz, which is in agreement with the 3 respiration cycles present in the 10 s capture.

Figure 5.30: (a) Measured respiration displacement with the 300 GHz radar. (b) Spectral representation of the measured data.

5.3 Future Work

This section describes some of the future work and applications related to the alternative uses for the 300 GHz and 122 GHz radars. First, improvements and future research lines for the NDT applications presented in Section 5.1 are detailed. Afterwards, some upgrades to the vital sign monitoring procedure from Section 5.2 are proposed. Lastly, further alternative applications and their preliminary tests are presented.

5.3.1 Data Acquisition

For an industrial implementation of the data acquisition technique proposed in Section 5.1, the main priorities are in sensor cost and size reductions, as well as guide optimization (regarding material and bit design). The application could also benefit from upgrades to the measurement set-up used, with emphasis on a 2-D positioning table to allow full 3-D imaging of the OUTs. In this section, a summary of some of the possible future research lines for both radars is presented.

Additionally, some specific work can be described for the 122 GHz device, such as:

- **Development of a Handheld Sensor.** The MMIC-based sensor is already substantially smaller in size than the 300 GHz radar. The sensor's current PCB size ($15 \times 15 \text{ cm}^2$), its use during measurements with a simple photography tripod as stabilizer, and the ability of embedded microprocessors to sample and process its data, prove that it is a few steps away from a handheld scanner. Further development is already being made in reducing the size of the PCB to just $5 \times 5 \text{ cm}^2$, and creating a 3-D printed handle.
- **MMIC Selection.** The MMICs tested during this thesis are not the only devices available from Silicon Radar [231]. Their research has recently produced additional chips, such as another 122 GHz MMIC with up to 20 GHz of bandwidth (thus allowing better range resolution) and a 245 GHz MIMO chip, which could provide an increase in performance. In this line, the MMIC with 20 GHz has already been acquired, fitted into an in-house made PCB, and tested with multiple measurements.
- **Multi-frequency Testing.** The sensor could be modified to incorporate a second already-available MMIC at 24 GHz also acquired from Silicon Radar. This chip presents a similar pin-out and RF architecture to the 122 GHz MMIC, allowing its use with a PCB slightly modified from the one already developed. With the two MMICs simultaneously scanning, multi-frequency evaluation of a material could be done, allowing the extraction of its parameters at different bands.

In order to finish a first prototype of the data acquisition sensor a processing algorithm to detect peaks in the spectrum and translate them to bits must be designed. For an initial development, high complexity for the algorithm is not expected, since the steps up to and including peak detection (already performed for the measurements presented) consist mostly on FFT processing. During peak detection, an adaptive threshold technique should be used in order to separate bit peaks from noise when illuminating the guide from different ranges and angles, where received power varies. A sliding window approach could be applied along the X dimension to smooth the results of peak selection and filter outliers. Afterwards, the length of the bits should be measured and compared against the design's baseline dimensions (with an error adjustment to compensate the increase from spot size). A block diagram of the proposed stages is shown in Fig. 5.31. Due to the small computational power required for these steps, the complete processing chain could be implemented in an embedded system, such as an FPGA, DSP, or some GPUs, to limit the size impact in the sensor.

Figure 5.31: Block diagram of the proposed bit-reading algorithm for the data acquisition technique.

In addition to algorithm development, a revision of the HDPE guide designs considered can prove beneficial for future implementations. Regarding alternative materials for the guides, multiple plastics aside from HDPE, as well as other dielectrics, have been previously evaluated in the mm-wave band [395, 396]. As for bit design, the triangular and rectangular shapes tested in Section 5.1.2 have been chosen for their similarity to corner reflectors often used as calibration targets in conventional radar systems, which provide powerful reflections due to large RCS. Nonetheless, as has been demonstrated through measurements, received power in the mm-wave band is strongly dependent on incidence angle, for which faceted surfaces provide varying results. In this regard, the study of smooth shapes such as cylindrical guides or spherical bits could provide improved designs resistant to sensor movement. Further additions to the guides, such as hydrophobic painting to disperse surface-accumulated water, can also be studied.

In relation to guide redesign, an expanded analysis of the bit alphabets used is necessary. The radar's ability to detect peaks up to tens of meters away allows for alphabets with bigger dimensionality than a binary code, by vertically stacking multiple cavities in a single column, as shown in Fig. 5.32. Additionally, a test with multiple guides joined in the Z dimension should be performed, to validate the possibility of simultaneous guide scanning. The development of upgraded alphabets can also be improved with the implementation of error-correction codes in the guide, further facilitating the accurate extraction of the information recorded in the guide.

Figure 5.32: Examples of information-density improvements for the data-acquisition technique. (a) shows the 2-bit-level single guide scheme, (b) shows a guide with increased bits in each X step, while (c) shows two guides stacked in the Z dimension.

Information density in the guide can be increased with the reduction of the radar's XY spot size and Z resolution. Although the latter depends primarily on the bandwidth transmitted, improvements to the windowing stage applied to the received signal could also provide better resolution. A more straightforward upgrade can be achieved in the cross-range axes,

through a custom design for the focusing lenses [397]. This design could even be expanded into an integrated lens antenna [398], eliminating the space between lenses and antennas, further reducing the dimensions of the sensor.

For experiments performed further along in the future, three main topics can be described:

- To validate the scanning capabilities images generated with continuous movement must be tested, instead of the step-by-step approach presented in previous sections. In the vehicle-mounted application of the technique a stable trajectory should first be used. This set-up holds a direct correlation to rail-guided vehicles, such as trains or subways, which follow a fixed route through infrastructure already established. In this regard, the application provides a solution to some of the precision and security problems of current train-positioning techniques [399, 400], such as Eurobalises [401], through placement of the guides in specific railway locations. Additionally, the radar could be jointly used with an additional positioning sensor in a fusion-based approach [402]. The reduced lateral movement of rail vehicles compared to cars, which produces less sensor misalignment and image distortion, allows for a simpler transition into an industrial prototype.
- Although it has been assumed for the guides to be planted at surface level, Ground-Penetrating Radar (GPR) systems offering underground images are common [403]. Although these are typically implemented at lower frequencies, benefitting from better penetration, some examples at the mm-wave exist [404]. This variation would provide the possibility of placing the guide below ground, reducing its exposure to adverse weather conditions and tear, and increasing privacy and security.
- Additional studies regarding the reflections from nearby materials around the guide are needed. For instance, in a pallet-tagging application the guide may be glued to a wooden crate, while for vehicle positioning the guide may be surrounded by various types of soil. Any time the sensor is misaligned reflections from these materials must be identified or suppressed to avoid reading errors.

5.3.2 Vital Sign Monitoring

Several future research lines and experiments regarding vital sign monitoring can be pursued with the groundwork presented in this thesis. The final objective would be to develop a prototype with diagnosis capabilities for a controlled clinical set-up. To achieve this goal, some useful improvements would be the following:

- **Simultaneous respiration and heart rate detection.** For the heartbeat-detection experiments in this thesis, the subject has been asked to hold his breath during the measurement. This is done to avoid masking the pulse-caused movements with those due to respiration, which present larger amplitude. Nonetheless, plenty radar systems with simultaneous respiration and heartbeat detection have been presented, using advanced signal processing strategies to separate the information from both signals [405–409].
- **Body movement cancellation.** All of the measurement scenarios presented here have used a static patient. Furthermore, the patient has been instructed to stand as still as possible, to avoid interference from body movements in the returned phase. Because movements related to vital signs have very low amplitudes, almost any movement by the subject will mask the information from both respiration and heart rate. However, in a real-life scenario, it cannot always be expected for the subject to be perfectly still. Although this is one of the main challenges affecting the technology [410], several methods have been proposed in related literature to deal with signal artifacts produced

by body movement [411–413]. To offer truly feasible prototypes, these methods must be included in the 122 GHz and 300 GHz radars.

- Measurement-verification support sensors.** For a complete vital sign monitoring system, support sensors must be used to validate the respiration and heart rate measurements of the radar. In the presented experiments, the respiration rate has been confirmed visually, while the heart rate has been monitored with an evaluation board. However, commercial sensors with increased functionalities, such as USB or WiFi communications for data logging and storage, are available. If, like the radar, the support sensors used are also non-contact, this provides an additional step forward into the future remote monitoring system. For this end, a Microsoft Kinect device can be used to remotely monitor both respiration and heart rate, and offer a secondary measurement to corroborate the radar's data. As an example, a first-generation Microsoft Kinect device has been used to track a subject's breathing remotely through the use of its directional microphones [414], as shown in Fig. 5.33. A reference trace corresponding to a Kinect capture from a silent scene is included for comparison. As can be seen, the starts of the recorded sound spikes match the instants where the radar sinusoid starts to rise, corresponding to the moment the subject exhales (when his chest deflates and he produces a hearable sound). Although already demonstrated in the literature, no heart rate detection with the Kinect has yet been performed during this thesis, because for this measurement the second-generation Kinect device is required. This is because this measurement uses slight changes in a subject's face depth and color (due to the flow of blood during pumping) to register the heart rate [415], for which the second-generation device's face tracking technology is necessary.

Figure 5.33: Joint measurement of a resting male subject's breathing rate using the 122 GHz radar and a Microsoft Kinect device.

- Multi-view measurements.** The low-cost of the 122 GHz radar, or the already-available multistatic configuration of the 300 GHz system, facilitate future multi-view vital sign monitoring measurements. In this scope, multiple synchronized sensors are used to measure a patient from different points of view. This can be useful to isolate body areas where a specific vital sign is more powerful, corroborate other sensors' measurements, or assign the detection individual vital sign parameters to each node. Such a multi-view architecture for this application has already been presented [416], with some examples reaching the THz range [417].
- Simultaneous imaging and monitoring.** The 300 GHz radar's ability to generate 3-D images of the patient can also be exploited. Measurements generating 3-D images

of the subject with simultaneous vital sign monitoring are possible with a fast-enough scanning and processing. In a series of continuous images, the pixels in each image can be treated as individual vital sign samples. After sufficient images, the chest and subcutaneous movements could be tracked and extracted to form a slow-time representation, for each FoV spot. In addition, the images themselves can be localized to specific body areas in order to offer more medical information, as done in burn injury evaluation [68], breast cancer analysis [67], or cornea imaging [69].

In addition to these main areas, the work presented here can be complemented with secondary improvements. In this regard, the effect of including more clothing during measurements should be analyzed to offer a more realistic monitoring scenario. A comparison among measurements where the radar aims at different body areas should also be performed, to establish the ideal location for each desired vital sign. Finally, a more extensive measurement campaign, encompassing subjects of varying age, gender, and body type [410] should be designed and performed.

5.3.3 Future Applications for Millimeter-wave Radars

In addition to the data acquisition and vital sign monitoring activities described in the previous sections, the following uses for millimeter-wave radars (particularly integrated sensors such as the 122 GHz MMIC-based device in this thesis) could also provide new applications.

5.3.3.1 Vibration Detection

An application enabled by the excellent performance of the 122 GHz radar in phase stability is the measurement of an object's vibrations. Because vibration movements usually have small amplitudes, cm-level range resolutions are not sufficient to accurately detect and track them. For this purpose, the μm -level phase stability of the 122 GHz radar is extremely useful. To test out the possibility of this application, several initial experiments have been performed.

The first experiment is based on targeting a slightly-open metallic closet door. During the 10-second measurement, the radar is continuously tracking the phase of the door's surface. At some point during the measurement, the door is suddenly closed, creating a sudden movement and vibration in the metallic material. The distance fluctuations associated to the returned phase [according to (3.6)] may be seen in Fig. 5.34. The exact instant where the door is closed can be perfectly recognized in the resulting graph, due to the large spike in distance fluctuation, followed by an oscillating aftermath due to vibrations. After applying an FFT to the data in Fig. 5.34(a), the different frequency components of the oscillation can be detected, which could possibly lead to registration of a material's resonant frequency. The radar's ability to detect the door's closing or opening is useful to home security applications, such as intruder detection (where the radar could be placed aiming at an entrance or window to register trespassings) [418].

This measurement scenario has been repeated, with the radar aiming at the floor near a human subject. During this continuous measurement, the subject abruptly falls to the ground, away from (but near) the radar's illuminated spot. The distance fluctuations are once again registered, and this time presented in Fig. 5.35. In this scenario, the specific spike where the fall happens can also be detected. The radar's capability to detect nearby impacts in the ground could be, for instance, useful in medical applications for the monitoring of disabled or senior patients, to detect sudden faintings, seizures, and other types of falls [419]. Because the damping factor of the floor is greater than the metallic door's, little oscillation is registered after the fall spike.

Figure 5.34: (a) Distance fluctuation from phase variations with the 122 GHz aiming at a metallic door that suddenly closes, and (b) spectral representation of the captured data.

Figure 5.35: (a) Distance fluctuation from phase variations with the 122 GHz aiming at the floor, near to where a human subject abruptly falls to the floor, and (b) spectral representation of the captured data.

Also within the vibration detection applications, a third set of measurements has been performed with the 122 GHz radar aiming at the underside of a thin wooden table [as shown in Fig. 5.36(a)]. During the continuous measurements, different objects are placed and removed from the table, directly above the radar. Because the table is narrow and made of a relatively flexible material, some bending is experienced when a heavy object is placed in it. In this regard, the radar is capable of measuring these μm deformations, thereby detecting the placement of an object, or even possibly the exact weight of the object placed (through the amplitude of the table's deformation). These types of measurements could be useful for domotic applications or non-contact weighting. A set of measurements, for different objects, is presented in Fig. 5.36(b). As can be seen from the results, increasing weights produce greater deformations. In all cases, once the object is removed from the table, the distance goes back to its original position.

Figure 5.36: Distance fluctuation from phase variations with the 122 GHz aiming at the underside of a table as several objects are placed and removed directly on top.

The final experiment deals with the measurement of the vibrations of a speaker when playing sound from different sinusoidal signals. The radar is placed aiming at the speaker's cone, which vibrates according to the frequency of the sounds it plays. By tracking the displacement of the speaker and performing an FFT on the resulting data, the frequency of the sound being played may be extracted. These types of measurements could be useful for long-range microphone systems [420], or even help diagnose speech patterns and impediments through the tracking of displacements in the larynx and rest of the vocal tract [421]. For these measurements, signals of various frequencies have been selected to demonstrate the wide range of detection. The resulting spectrums are presented in Fig. 5.37. Notice that the measurement at 100 Hz presents various harmonics in the spectrum. This is because for this sound, the signal's amplitude was deliberately set too high and the speaker is forced to saturate its value, thus playing the sound of a pulse train instead of a sinusoid.

Additional applications for which a vibration-detecting sensor is needed are endless. For instance, in structural engineering, such a sensor could be used to measure the resonation caused in bridges, buildings, or tunnels as vehicles circulate nearby [422]. In human-machine interfaces, an FMCW radar chip from Google has recently been used to recognize different hand gestures through their motion [75]. The detection of micro-Doppler signals can also be useful for other biomedical and security applications, such as human gait analysis and recognition [423].

Figure 5.37: Spectral representation of the distance fluctuation from phase variations with the 122 GHz aiming at a speaker as sinusoidal signals at (a) 10 Hz, (b) 100 Hz, and (c) 2000 Hz are played.

5.3.3.2 Multi-node Sensor Network

A promising application for the developed 122 GHz radar is the implementation of a multi-node sensor network. Such a network would be composed of multiple 122 GHz radars working jointly or independently to measure one or several targets. The use of several sensors at once would allow measurements from different points of view. In the case of range measurements, a sensor network with three or more nodes (whose position is perfectly known) could be used to triangulate the exact location of a target in the scene. Conversely, instead of measuring the same parameter, each node could be tasked with the measurement of a different magnitude, such as range, Doppler velocity, or direction of arrival. This architecture is only plausible thanks to the low-cost and reduced-size of the developed radars. Additionally, the original 122 GHz commercial kit incorporated WiFi communication for its processed digital samples, which could also be implemented in the modified version to implement a fully-wireless distributed network.

Such a multi-node network has already been evaluated in the literature with a variation of the 122 GHz radar's commercial version [424–427]. In these works, distributed two-node and three-node configurations has been used to recognize the contour of a static object, as well as offer better location detection through triangulation. In addition, both monostatic and bistatic distances have been assessed, concluding that the use of bistatic measurements offers an increase in ranging performance.

With these previous results in mind, a proposed architecture is illustrated in Fig. 5.38, comparing supervised (where the nodes are slaved to a master controller) versus unsupervised (where each node operates independently) configurations. In the unsupervised scheme, each node operates independently, and it is up to the later signal processing to separate each node's information. In the supervised version, the nodes are synchronized and controlled by the PC, through the use of a master board. The master board distributes a universal clock signal to the nodes and ADC, for system coherence and to reduce phase noise after decorrelation in bistatic measurements. The master board also distributes a trigger signal to initiate each node's transmission at the same time, to personalize a specific time offset for each node, or to set-up a time-multiplex scheme. Finally, the master board configures the transmitted waveform for each node, which can also be equal for all sensors or node-specific. This could be used, for instance, to simultaneously scan the target with both CW and FMCW nodes, using different bandwidths and center frequencies.

The specific applications of these networks are abundant. The multi-view aspect of the network would of course be useful in security applications [429], such as perimeter surveillance, but also for medical applications [428], such as vital sign monitoring from various

angles or for various biomedical signals. Furthermore, the ability to measure Doppler velocity from various angles would be extremely advantageous in traffic surveillance scenarios. The mounting of the individual nodes in unmanned aerial vehicles could also lead to multiple applications based on SAR imaging [430], such as crop monitoring, urban land analysis, or geological evaluations.

Figure 5.38: Block diagrams comparing multi-node architectures in (a) unsupervised and (b) supervised settings.

With this multistatic architecture in mind, a new 122 GHz radar is currently under development. The new sensor has reduced dimensions, increased bandwidth, and additional synchronization features (such as outputs for clock and trigger distribution). In addition, FMCW interference and multi-node synchronization for the 122 GHz radars has already been tested for the imaging array configuration presented in Section 3.2.5. However, many undefined variables still remain, such as optimal node positioning, compensation of non-linearities caused by mixing chirps from different synthesizers, or efficient sampling techniques.

Chapter 6

Conclusions and Future Work

6.1 Conclusions

This thesis has been primarily focused on improving the imaging performance of security screening radars in the mm- and submm-wave bands. To achieve this goal, the work performed has been centered on two radar systems, at 300 GHz and 122 GHz, respectively. Aside from an initial introduction, the thesis has been divided into two first chapters (one for each system), with a third chapter describing algorithms for the automatic detection of threat relevant to both radars, and an additional fourth chapter detailing the radars' use in other applications aside from security imaging.

The first part of the thesis is focused on a pre-existing 300 GHz standoff imaging radar for the detection of person-borne concealed weapons. The imager has been modified, at hardware and software levels, to improve several key areas. In particular, a real-time signal and image processing architecture based on parallel computing has been implemented, achieving a ~ 10 fps state-of-the-art image refresh speed, crucial for accurate movement representation. Additionally, a multistatic imaging architecture based on multi-bistatic measurements around a target has been proposed and tested, demonstrating the advantages of multi-view imaging to improve the effective FoV under specular reflection effects.

The second part of the thesis details the development of a hybrid imager using a modified 122 GHz commercial radar and a COTS 94 GHz passive RX, as well as an optical sensor. To establish the hybrid system MMIC-based commercial products have been customized into FMCW sensors capable of being connected into imaging arrays or sensor networks. The modified radar has demonstrated excellent performance regarding chirp linearity, power ripple, and phase stability, among others. The active and passive images are combined using custom image fusion code, using a COTS optical sensor to filter unwanted outlier noise. The hybrid results obtained show that the combined and simultaneous use of multiple sensor types allows detection of threats in scenarios where single-sensor imagers would fail.

The third part of the thesis presents the design and evaluation of ATD algorithms for both active and passive images. In the case of the active images, the dataset was a large publicly-available collection created with the use of a mm-wave radar. For the passive images, an in-house dataset was created with the use of a COTS IR camera. Individual models were developed for each type of sensor, based on an AlexNet CNN architecture pre-trained on the ImageNet dataset. Results show the feasibility of using convolutional neural networks to detect concealed threats in both cases with accuracies higher than 99% and training times of just ~ 30 minutes per model.

The fourth, and final, part of the thesis describes the use of both radar systems in two applications aside from security imaging. The first application is a data acquisition technique from pre-recorded low-cost plastic guides using FMCW imaging, previously-proposed for vehicle positioning, but adaptable to multiple industrial fields. The second application tested has been the remote detection and monitoring of a human's vital signs. In this regard, the radars have been able to capture a subject's heartbeat and respiration under various scenarios involving different standoff ranges and breathing rates. Both applications are extensively detailed in Chapter 5.

6.2 Future Work

Throughout all sections of this thesis several future research opportunities have been presented, which are all valid approaches that should be addressed. Nonetheless, the main objective for any future work in relation to this thesis should be the development of a prototype imaging system that integrates all the improvements described in these chapters and can perform in a true security environment. Ideally, this prototype would comply with the requirements of stakeholders such as the TSA, which call for checkpoints with high passenger throughput and automatic screening technologies. To this end, the rise of COTS MMIC-based TRXs and the work presented in Chapter 3 provide a great starting point for a cost-efficient implementation. For a future prototype, three main areas of research can be defined:

- **Hybrid Imaging.** As has been detailed in Chapter 3, hybrid imaging is one of the recent trends in state-of-the-art security imagers to improve the detection capabilities in checkpoints. This improvement has been demonstrated with various imaging scenarios in this thesis. A future prototype imager will need to include multiple types of sensors (millimeter-wave active, millimeter-wave passive, IR, optical, etc...) to account for any kind of detection scenario. In this regard, this thesis has already shown the capabilities of an imager combining millimeter-wave active and passive data independently with optical information. Any further work should start with developing an image fusion algorithm to combine all three types of images in real-time, designing a focusing scheme adapted to multi-type sensors in an array configuration, and analyzing the addition of sensors working at different frequencies such as IR.
- **Multistatic Imaging.** Similarly, multistatic imaging is the second current trend in security imagers that aims to improve their performance. As explained and demonstrated in Chapter 2, multistatic imaging is necessary when working in a real-life environment where targets can move freely through the checkpoint. Additionally, if done in a stand-off configuration, further advantages in equipment security and subject compliance are gained. To this end, future imagers will need to include multiple TRXs distributed around the checkpoint's walkway. The future work in this topic will need to start with study and simulation of the optimal placement and number of TRXs in a given scene, as well as designing the appropriate signal or image combination strategy.
- **Automatic Threat Detection.** The evolution of threat detection strategies from visual and manual confirmation to automatic localization algorithms is the third trend in modern imagers. As described in Chapter 4, in order to reduce bottlenecks and ensure constant detection rates, automatic threat detection algorithms based on machine learning techniques are a necessity. The results in this thesis show the feasibility of adapting deep learning models trained on visible-spectrum images to detect concealed threats in a limited set of millimeter-wave imaging scenarios. For a future prototype imager, the first step should be increasing the scenarios included in the training dataset, followed by the design of a custom CNN architecture specific for millimeter-wave images.

Bibliography

- [1] J. C. Wiltse. "History of millimeter and submillimeter waves". *IEEE Trans. Microw. Theory Tech.*, vol. 32, no. 9, pp. 1118-1127, Sept. 1984.
- [2] P. De Maagt, P. H. Bolivar, & C. Mann. "Terahertz science, engineering and systems - From space to earth applications". *Encyclopedia of RF and Microwave Engineering*, Apr. 2005.
- [3] J. Ramsey. "Millimeter wave research in the 1890's". *IEEE Spectr.*, vol. 4, no. 5, Dec. 1967.
- [4] E. D. Palik. "History of far-infrared research. I. The Rubens era". *J. Opt. Society America*, vol. 67, no. 7, pp. 857-865, Jul. 1977.
- [5] M. I. Skolnik. *Radar handbook*. McGraw-Hill, 1990.
- [6] W. Gordy. "Millimeter and submillimeter waves in physics." *Proc. Symp. Millimeter Waves*, vol. 9, pp. 1 - 23, Apr. 1959.
- [7] A. B. Crawford, & D. C. Hogg. "Measurement of atmospheric attenuation at millimeter wavelengths". *Bell Syst. Tech. J.*, vol. 35, no. 4, pp. 907-916, 1956.
- [8] J. C. Wiltse, S. P. Schlesinger, & C. M. Johnson. "Back-scattering characteristics of the sea in the region from 10 to 50 KMC". *Proc. IRE*, vol. 45, no. 2, pp. 220-228, 1957.
- [9] J. C. Wiltse. "Some characteristics of dielectric image lines at millimeter wavelengths". *IRE Trans. Microw. Theory Tech.*, vol. 7, no. 1, pp. 65-69, Jan. 1959.
- [10] C. M. Johnson, D. M. Slager, & D. D. King. "Millimeter waves from harmonic generators". *Review of Scientific Instruments*, vol. 25, no. 3, pp. 213-217, Mar. 1954.
- [11] Meinel, Holger H. "Commercial applications of millimeterwaves: history, present status, and future trends." *IEEE Trans. Microw. Theory Tech.*, vol. 43, no. 7, pp. 1639-1653, Jul. 1995.
- [12] P. H. Siegel. "THz instruments for space". *IEEE Trans. Antennas Propag.*, vol. 55, no. 11, pp. 2957-2965, Nov. 2007.
- [13] M. F. Kessler, et al. "The Infrared Space Observatory (ISO) mission". *Astron. Astrophys*, vol. 315, L27-L31, 1996.
- [14] G. J. Melnick, et al. "The submillimeter wave astronomy satellite: Science objectives and instrument description". *The Astrophysical Journal Letters*, vol. 539, no.2, L77, 2000.

- [15] I. D. Robertson, & S. Lucyszyn. *RFIC and MMIC Design and Technology*. IET, 2009.
- [16] P. H. Siegel. "Terahertz technology in biology and medicine". *IEEE Trans. Microw. Theory Tech.*, vol. 52, no. 10, pp. 2438-2447, Oct. 2004.
- [17] C. M. Armstrong. "The truth about terahertz". *IEEE Spectr.*, vol. 49, no. 9, pp. 36-41, Sept. 2012.
- [18] P. H. Siegel. "Terahertz technology". *IEEE Trans. Microw. Theory Tech.*, vol. 50, no. 3, pp. 910-928, Mar. 2002.
- [19] J. A. Nanzer. "A review of microwave wireless techniques for human presence detection and classification". *IEEE Trans. on Microw. Theory Tech.*, vol. 65, no. 5, pp. 1780-1794, May 2017.
- [20] R. Appleby, & H. B. Wallace. "Standoff detection of weapons and contraband in the 100 GHz to 1 THz region". *IEEE Trans. Antennas Propag.*, vol. 55, no. 11, pp. 2944-295, Nov. 2007.
- [21] N. C. Currie, & C. E. Brown, *Principles and applications of millimeter-wave radar*. Artech House, 1987
- [22] G. Ortiz-Jiménez, et al. "Simulation framework for a 3-D high-resolution imaging radar at 300 GHz with a scattering model based on rendering techniques". *IEEE Trans. THz Sci. Technol.*, vol. 7, no. 4, pp. 404-414, July 2017.
- [23] R. Knipper, et al. "THz absorption in fabric and its impact on body scanning for security application". *IEEE Trans. THz Sci. Technol.*, vol. 5, no. 6, pp. 999-1004, Nov. 2015.
- [24] H.-R. Chen. *FMCW radar jamming techniques and analysis*. Naval Postgraduate School, 2013.
- [25] E. G. Njoku. "Passive microwave remote sensing of the earth from space - A review". *Proc. IEEE*, vol. 70, no. 7, pp. 728-750, Jul. 1982.
- [26] A. Wootten, & A. R. Thompson. "The Atacama large millimeter/submillimeter array". *Proc. IEEE*, vol. 97, no. 8, pp. 1463-1471, Aug. 2009.
- [27] G. L. Pilbratt, et al. "Herschel Space Observatory-An ESA facility for far-infrared and submillimetre astronomy". *Astronomy & Astrophysics*, vol. 518, L1, 2010.
- [28] K.-H. Glassmeier, et al. "The Rosetta mission: flying towards the origin of the solar system". *Space Science Reviews*, vol. 128, no. 1-4, pp. 1-21, Feb. 2007.
- [29] J. A. Tauber, et al. "Planck pre-launch status: The Planck mission". *Astronomy & Astrophysics*, vol. 520, A1, 2010.
- [30] D. P. Woody, et al. "CARMA: a new heterogeneous millimeter-wave interferometer". *Millimeter and Submillimeter Detectors for Astronomy II*, vol. 5498, pp. 30-42, Oct. 2004.
- [31] K. D. Klaes, et al. "An introduction to the EUMETSAT polar system". *Bulletin of the American Meteorological Society*, vol. 88, no. 7, pp. 1085-1096, Jul. 2007.
- [32] E. T. Young, et al. "Early science with SOFIA, the stratospheric observatory for infrared astronomy". *The Astrophysical Journal Letters*, vol. 749, no. 2, L17, Mar. 2012.
- [33] K.-I. Kikuchi, et al. "Overview and early results of the Superconducting Submillimeter-Wave Limb-Emission Sounder (SMILES)". *Journal of Geophysical Research: Atmospheres*, vol. 115, D23, Dec. 2010.

- [34] P. T. Ho, J. M. Moran, & K. Yung Lo. "The submillimeter array". *The Astrophysical Journal Letters*, vol. 616, no. 1, L1, Oct. 2004.
- [35] H. J. Hartfuss, T. Geist, & M. Hirsch. "Heterodyne methods in millimetre wave plasma diagnostics with applications to ECE, interferometry and reflectometry". *Plasma Physics and Controlled Fusion*, vol. 39, no. 11, pp. 1693, Nov. 1997.
- [36] X. P. Lu, & M. Laroussi. "Electron density and temperature measurement of an atmospheric pressure plasma by millimeter wave interferometer". *Applied Physics Letters*, vol. 92, no. 5, 051501, Feb. 2008.
- [37] Y. Wang, et al. "Millimeter-wave imaging of magnetic fusion plasmas: technology innovations advancing physics understanding". *Nuclear Fusion*, vol. 57, no. 7, 072007, Mar. 2017.
- [38] D. W. Van Der Weide, J. Murakowski, & F. Keilmann. "Gas-absorption spectroscopy with electronic terahertz techniques". *IEEE Trans. Microw. Theory Tech.*, vol. 48, no. 4, pp. 740-743, Apr. 2000.
- [39] N. Gopalsami, S. Bakhtiari, T. W. Elmer II, & A. C. Raptis. "Application of millimeter-wave radiometry for remote chemical detection". *IEEE Trans. Microw. Theory Tech.*, vol. 56, no. 3, pp. 700-709, Mar. 2008.
- [40] M. R. Leahy-Hoppa, M. J. Fitch, & R. Oslander. "Terahertz spectroscopy techniques for explosives detection". *Analytical and bioanalytical chemistry*, vol. 395, no. 2, pp. 247-257, Sept. 2009.
- [41] K. B. Cooper, et al. "Atmospheric humidity sounding using differential absorption radar near 183 GHz". *IEEE Geosci. Remote Sens. Lett.*, vol. 15, no. 2, pp. 163-167, Feb. 2018.
- [42] G. Rubio-Cidre, et al. "Understanding cloud dynamics using a ground-based radar at 94 GHz". *2017 IEEE Radar Conference (RadarConf)*, pp. 1137-1141, May 2017.
- [43] D. L. Wu, J. H. Jiang, & C. P. Davis. "EOS MLS cloud ice measurements and cloudy-sky radiative transfer model". *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 5, pp. 1156-1165, May 2006.
- [44] C. Surussavadee, & D. H. Staelin. "Global millimeter-wave precipitation retrievals trained with a cloud-resolving numerical weather prediction model, Part I: Retrieval design". *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 1, pp. 99-108, Jan. 2008.
- [45] W. J. Wilson, & P. R. Schwartz. "Diurnal variations of mesospheric ozone using millimeter-wave measurements". *Journal of Geophysical Research: Oceans*, vol. 86, no. C8, pp. 7385-7388, Aug. 1981.
- [46] C. L. Croskey, et al. "The Millimeter Wave Atmospheric Sounder (MAS): A shuttle-based remote sensing experiment". *IEEE Trans. Microw. Theory Tech.*, vol. 40, no. 6, pp. 1090-1100, Jun. 1992.
- [47] S. Ayhan, et al. "Millimeter-wave radar sensor for snow height measurements". *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 2, pp. 854-861, Feb. 2017.
- [48] E. R. Brown, et al. "Wide-band radiometry for remote sensing of oil films on water". *IEEE Trans. Microw. Theory Tech.*, vol. 46, no. 12, pp. 1989-1996, Dec. 1998.
- [49] F. J. Wentz. "Measurement of oceanic wind vector using satellite microwave radiometers". *IEEE Trans. Geosci. Remote Sens.*, vol. 30, no. 5, pp. 960-972, Sept. 1992.

- [50] D. M. Sheen, D. L. McMakin, & T. E. Hall, "Three-dimensional millimeter-wave imaging for concealed weapon detection", *IEEE Trans. Microw. Theory Tech.*, vol. 49, no. 9, pp. 1581 - 1592, Sept. 2001.
- [51] D. L. McMakin, D. M. Sheen, J. W. Griffin, & W. M. Lechelt. "Extremely high-frequency holographic radar imaging of personnel and mail". *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense V*, vol. 6201, p. 62011W, May 2006.
- [52] T. S. Rappaport, et al. "Millimeter wave mobile communications for 5G cellular: It will work!". *IEEE Access*, vol. 1, pp. 335-349, May 2013.
- [53] T. Nitsche, et al. "IEEE 802.11 ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi". *IEEE Commun. Mag.*, vol. 52, no. 12, pp. 132-141, Dec. 2014.
- [54] S. L. Cotton, W. G. Scanion, & B. K. Madahar. "Millimeter-wave soldier-to-soldier communications for covert battlefield operations". *IEEE Commun. Mag.*, vol. 47, no. 10, pp. 72-81, Oct. 2009.
- [55] J. H. Wehling. "Multifunction millimeter-wave systems for armored vehicle application". *IEEE Trans. Microw. Theory Tech.*, vol. 53, no. 3, pp. 1021-1025, Mar. 2005.
- [56] J. N. Pelton. "New millimeter, terahertz, and light-wave frequencies for satellite communications". *Handbook of satellite applications*, pp. 413-429, 2017.
- [57] P. J Shull. *Nondestructive evaluation: theory, techniques, and applications*. CRC press, 2016.
- [58] Y. L. Hor, J. F. Federici, & R. L. Wample. "Nondestructive evaluation of cork enclosures using terahertz/millimeter wave spectroscopy and imaging". *Applied Optics*, vol. 47, no. 1, pp. 72-78, Jan. 2008.
- [59] Y. Yashchyshyn, & K. Godziszewski. "A new method for dielectric characterization in sub-THz frequency range". *IEEE Trans. THz Sci. Technol.*, vol. 8, no. 1, pp. 19-26, Jan. 2018.
- [60] W. Fan, et al. "Dielectric properties of coal in the terahertz frequency region of 100-500 GHz". *Fuel*, vol. 188, pp. 246-253, Dec. 2015.
- [61] J. Dash, et al. "Non-destructive inspection of internal defects in concrete using continuous wave 2D terahertz imaging system". *2013 38th International Conference on Infrared, Millimeter, and Terahertz Waves (IRMMW-THz)*, pp. 1-2, Sept. 2013.
- [62] K. Kawase, et al. "Non-destructive terahertz imaging of illicit drugs using spectral fingerprints". *Optics express*, vol. 11, no. 20, pp. 2549-2554, Oct. 2003.
- [63] A. Abina, et al. "Applications of terahertz spectroscopy in the field of construction and building materials". *Applied spectroscopy reviews*, vol. 50, no. 4, pp. 279-303, Apr. 2015.
- [64] Z. Yan, et al. "Research progress of terahertz wave technology in food inspection". *Terahertz Physics, Devices, and Systems*, vol. 6373, pp. 63730R, Oct. 2006.
- [65] J. C. Lin. "Microwave sensing of physiological movement and volume change: A review". *Bioelectromagnetics*, vol. 13, no. 6, pp. 557-565, 1992.
- [66] C. Li, et al. "A review on recent advances in Doppler radar sensors for noncontact healthcare monitoring". *IEEE Trans. Microw. Theory Tech.*, vol. 61, no. 5, pp. 2046-2060, May 2013.

- [67] M. Klemm, et al. "Microwave radar-based differential breast cancer imaging: Imaging in homogeneous breast phantoms and low contrast scenarios". *IEEE Trans. Antennas Propag.*, vol. 58, no. 7, pp. 2337-2344, Jul. 2010.
- [68] Y. Gao, & R. Zoughi. "Millimeter wave reflectometry and imaging for noninvasive diagnosis of skin burn injuries". *IEEE Trans. Instrum. Meas.*, vol. 66, no. 1, pp. 77-84, Jan. 2017.
- [69] S. Sung, et al. "THz imaging system for in vivo human cornea". *IEEE Trans. THz Sci. Technol.*, vol. 8, no. 1, pp. 27-37, Jan. 2018.
- [70] I. Skog, & P. Handel. "In-car positioning and navigation technologies-A survey". *IEEE Trans. Intell. Transp. Syst.*, vol. 10, no. 1, pp. 4-21, Mar. 2009.
- [71] I. Gresham, et al. "Ultra-wideband radar sensors for short-range vehicular applications". *IEEE Trans. Microw. Theory Tech.*, vol. 52, no. 9, pp. 2105-2122, Sept. 2004.
- [72] T.-Y. Lee, et al. "Application of W-band FMCW radar for road curvature estimation in poor visibility conditions". *IEEE Sensors J.*, vol. 18, no. 13, pp. 5300-5312, Jul. 2018.
- [73] D. Jasteh, et al. "Experimental low-terahertz radar image analysis for automotive terrain sensing". *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 4, pp. 490-494, Apr. 2016.
- [74] X. Yang, et al. "Three-dimensional imaging of space debris with space-based terahertz radar". *IEEE Sensors J.*, vol. 18, no. 3, pp. 1063-1072, Feb. 2018.
- [75] J. Lien, et al. "Soli: Ubiquitous gesture sensing with millimeter wave radar". *ACM Transactions on Graphics (TOG)*, vol. 35, no. 4, pp. 142, Jul. 2016.
- [76] M. Rangwala, F. Wang, & K. Sarabandi. "Study of millimeter-wave radar for helicopter assisted landing system". *2007 IEEE International Geoscience and Remote Sensing Symposium*, pp. 777-780, Jul. 2007.
- [77] T.J. Hewison. "Airborne measurements of forest and agricultural land surface emissivity at millimeter wavelengths". *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 2, pp. 393-400, Feb. 2001.
- [78] A. Kohmura, et al. "Optical fiber connected millimeter-wave radar for FOD detection on runway". *2013 European Radar Conference*, pp. 41-44, Oct. 2013.
- [79] G. Brooker, R. Hennessey, C. Lobsey, M. Bishop, & E. Widzyk-Capehart. "Seeing through dust and water vapor: Millimeter wave radar sensors for mining applications". *Journal of Field Robotics*, vol. 24, no. 7, pp. 527-557, Jul. 2007.
- [80] M. Jiao, et al. "A novel radar sensor for the non-contact detection of speech signals". *Sensors*, vol. 10, no. 5, pp. 4622-4633, 2010.
- [81] T. S. Rappaport, et al. "Overview of millimeter wave communications for fifth-generation (5G) wireless networks-with a focus on propagation models". *IEEE Trans. Antennas Propag.*, vol. 65, no. 12, pp. 6213-6230, Dec. 2017.
- [82] R. Appleby, D. A. Robertson, & D. Wikner. "Millimeter wave imaging: a historical review". *Passive and Active Millimeter-Wave Imaging XX*, vol. 10189, pp.1018902, May 2017.
- [83] Department of Homeland Security - Passive Millimeter Wave Detectors Market Survey Report [Online]: https://www.dhs.gov/sites/default/files/publications/PassiveMmWaveDetect-MSR_0614-508.pdf

- [84] Congressional Research Service - Airport Body Scanners: The Role of Advanced Imaging Technology in Airline Passenger Screening [Online]: <https://fas.org/sgp/crs/homesecc/R42750.pdf>
- [85] T. Hochrein. "Markets, availability, notice, and technical performance of terahertz systems: historic development, present, and trends". *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 36, no. 3, pp. 235-254, Mar. 2015.
- [86] S. S. Dhillon, et al. "The 2017 terahertz science and technology roadmap". *Journal of Physics D: Applied Physics*, vol. 50, no. 4, pp. 043001, Jan. 2017.
- [87] K.-E. Peiponen, A. Zeidler, & Makoto K.-G. *Terahertz spectroscopy and imaging*, Springer, 2012.
- [88] I. Minin. *Microwave and Millimeter Wave Technologies from Photonic Bandgap Devices to Antenna and Applications*, IntechOpen, 2010.
- [89] P. H. Siegel, & R. J. Dengler. "Terahertz heterodyne imaging part II: instruments". *International journal of infrared and millimeter waves*, vol. 27, no. 5, pp. 631-655, Apr. 2006.
- [90] R. Appleby, & R. N. Anderton. "Millimeter-wave and submillimeter-wave imaging for security and surveillance". *Proc. IEEE*, vol. 95, no. 8, pp. 1683-1690, Aug. 2007.
- [91] J. Hassel et al. "Dual-band submillimeter-wave kinetic inductance bolometers and an imaging system for contraband object detection". *Passive and Active Millimeter-Wave Imaging XXI*, vol. 10634, p. 106340F, May 2018.
- [92] D. A. Robertson, et al. "The CONSORTIS 16-channel 340 GHz security imaging radar". *Passive and Active Millimeter-Wave Imaging XXI*, vol. 10634, pp. 1063409, May 2018.
- [93] R. Herschel, S. A. Lang, & N. Pohl. "MIMO imaging for next generation passenger security systems". *Proceedings of EUSAR 2016: 11th European Conference on Synthetic Aperture Radar*, pp. 1 - 4, June 2016
- [94] T. Zwick, et al. "Pea-Sized mmW Transceivers: QFN-Based Packaging Concepts for Millimeter-Wave Transceivers". *IEEE Microw. Mag.*, vol. 18, no. 6, pp. 79 - 89, Aug. 2017.
- [95] F. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform". *Proc. IEEE*, vol. 66, pp. 51 - 83, 1978.
- [96] C. Pérez-Moreno, & J. Grajal. "Broadband frequency triplers for radar applications in the THz range". *29th IEEE International Symposium on Space THz Technology (ISSTT2018)*, pp. 210-211, Mar. 2018
- [97] G. Rubio-Cidre, et al. "DDS-based signal-generation architecture comparison for an imaging radar at 300 GHz". *IEEE Trans. Instrum. Meas.*, vol. 64, no. 11, pp. 3085-3098, Nov. 2015.
- [98] Department of Homeland Security (2018) - *Apex Screening at Speed* [Online]. At: https://www.dhs.gov/sites/default/files/publications/Apex%20Screening%20at%20Speed%20508%20Approved%20Final_0.pdf
- [99] Transport Security Administration (2018) - *Advanced Integrated Passenger and Baggage Screening Technologies* [Online]. At: <https://www.dhs.gov/sites/default/files/publications/TSA%20-%20Advanced%20Integrated%20Passenger%20Screening%20Technologies.pdf>

- [100] B. Kapilevich, & M. Einat. "Detecting hidden objects on human body using active millimeter wave sensor". *IEEE Sensors J.*, vol. 10, no. 11, pp. 1746-1752, Nov. 2010.
- [101] L3 ProVision 2 - L3 Security & Detection Systems [Online] <https://www.sds.l3t.com/advancedimaging/provision-at.htm>
- [102] J. Grajal, et al. "3-D high-resolution imaging radar at 300 GHz with enhanced FoV", *IEEE Trans. Microw. Theory Tech.*, vol. 63, no. 3, pp. 1097 - 1107, January 2015.
- [103] K. B. Cooper, et al. "Penetrating 3-D imaging at 4-and 25-m range using a submillimeter-wave radar", *IEEE Trans. Microw. Theory Tech.*, vol. 56, no. 12, pp. 2771-2778, Dec. 2008.
- [104] F. García-Rial, L. Úbeda-Medina, & J. Grajal, "Real-time GPU-based image processing for a 3-D THz radar". *IEEE Trans. Parallel Distrib. Syst.*, vol. 28, no. 10, pp. 2953 - 2964, Oct. 2017.
- [105] F. García-Rial, et al. "Combining commercially available active and passive sensors into a millimeter-wave imager for concealed weapon detection". *IEEE Trans. Microw. Theory Tech.*, vol. 67, no. 3, pp. 1167-1183, Mar. 2019.
- [106] O. Yurduseven, K. Cooper, & G. Chattopadhyay. "Point-spread-function (PSF) characterization of a 340 GHz imaging radar using acoustic levitation". *IEEE Trans. THz Sci. Technol.*, vol. 9, no. 1, pp. 20-26, Jan. 2019.
- [107] D. A. Robertson, et al. "A high frame rate, 340 GHz 3D imaging radar for security". *2018 IEEE Radar Conference (RadarConf18)*, pp. 55-60, Apr. 2018.
- [108] D. M. Sheen, et al. "Standoff concealed weapon detection using a 350-GHz radar imaging system". *Passive Millimeter-Wave Imaging Technology XIII*, vol. 7670, pp. 767008, Apr. 2010.
- [109] S. S. Ahmed, A. Genghammer, A. Schiessl, & L. Schmidt. "Fully electronic E-band personnel imager of 2 m² aperture based on a multistatic architecture". *IEEE Trans. Microw. Theory Tech.*, vol. 61, no. 1, pp. 651 - 657, Dec. 2012.
- [110] T. Bryllert, et al. "220-GHz imaging radar with 1-Hz frame rate using an array of homodyne transceivers". *Proc. of SPIE*, vol. 10634, May 2018.
- [111] S. Gui, J. Li, & Yiming Pi. "Security imaging for multi-target screening based on adaptive scene segmentation with terahertz radar". *IEEE Sensors J.*, vol. 19, no. 7, pp. 2675-2684, Apr. 2018.
- [112] S. Stanko, et al. "Millimeter resolution SAR imaging of infrastructure in the lower THz region using MIRANDA-300". *2016 46th European Microwave Conference (EuMC)*, pp. 1505-1508, Oct. 2016.
- [113] B. Cheng, et al. "340 GHz 3-D imaging radar with 4Tx-16Rx MIMO array". *IEEE Trans. THz Sci. Technol.*, vol. 8, no. 5, pp. 509-519, Sept. 2018.
- [114] J. Grzyb, K. Statnikov, N. Sarmah, B. Heinemann, & U. R. Pfeiffer. "A 210-270-GHz circularly polarized FMCW radar with a single-lens-coupled SiGe HBT chip". *IEEE Trans. THz Sci. Technol.*, vol. 6, no. 6, pp. 771-783, Nov. 2016.
- [115] A. Mostajeran, et al. "A high-resolution 220-GHz ultra-wideband fully integrated ISAR imaging system". *IEEE Trans. Microw. Theory Tech.*, vol. 67, no. 1, pp. 429-442, Jan. 2019.

- [116] B. M. Kent, T. Van, T. Van, K. Hamblin, & J. Westhoven. "Measurements of incident radio frequency power levels from the L3 technologies provision body scanner for the national academy of science". *2017 Antenna Measurement Techniques Association Symposium (AMTA)*, pp. 1-6, Oct. 2017.
- [117] Smiths Detections - eqo [Online]: <https://www.smithsdetection.com/products/eqo/>
- [118] ThruVision - ThruVision TS4-SC [Online]: <http://thruvision.com/wp-content/uploads/2018/11/1810-TV-TS4-SC-Datasheet-v1.3.pdf>
- [119] Asqella Oy. - ARGON [Online]: https://asqella.com/wp-content/uploads/2018/02/ARGON_brochure_web-1-copy.pdf
- [120] E. Heinz, et al. "Passive 350 GHz video imaging systems for security applications". *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 36, no. 10, pp. 879-895, Oct. 2015.
- [121] B. Mencía Oliva. *Desarrollo y aplicaciones de radares de alta resolución en bandas milimétricas*. PhD Thesis, ETSIT-UPM, 2014.
- [122] F. García-Rial, et al. "Optimization of a compact THz imaging radar for real-time operation". *Millimeter Waves (GSMM) & ESA Workshop on Millimetre-Wave Technology and Applications, 2016 Global Symposium on*, pp. 1-4, Jun. 2016.
- [123] J. Grajal, et al. "Compact radar front-end for an imaging radar at 300 GHz". *IEEE Trans. THz Sci. Technol.*, vol. 7, no. 3, May 2017.
- [124] J. Vankka, & K. A. I. Halonen. *Direct digital synthesizers: theory, design and applications*, Springer Science & Business Media, 2013.
- [125] A. Garcia-Pino, "Bifocal reflector antenna for a standoff radar imaging system with enhanced field of view". *IEEE Trans. Antennas Propag.*, vol. 62, no. 10, pp. 4997 - 5006, October 2014.
- [126] K. B. Cooper, et al. "THz imaging radar for standoff personnel screening". *IEEE Trans. THz Sci. Technol.*, vol. 1, no. 1, pp. 169-182, Sept. 2011.
- [127] G. Rubio Cidre, *Technological developments for millimeter-wave radars. Desarrollos tecnológicos para radares en bandas de milimétricas*. PhD Thesis, ETSIT-UPM, 2017.
- [128] K. B. Cooper & G. Chattopadhyay. "Submillimeter-wave radar: solid-state system design and applications". *IEEE Microw. Mag.*, vol. 15, no. 7, pp. 51 - 67, Dec. 2014.
- [129] D. A. Robertson, D. G. Macfarlane, & T. Bryllert. "220 GHz wideband 3D imaging radar for concealed object detection technology and phenomenology studies". *Passive and Active Millimeter-Wave Imaging XIX*, vol. 9830, pp. 983009, May 2016.
- [130] M. Caris, et al. "Very high resolution radar at 300 GHz". *2014 44th European Microwave Conference (EuMC)*, pp. 1797 - 1799, Oct. 2014
- [131] B. Kapilevich, et al. "330 GHz FMCW image sensor for homeland security applications". *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 31, no. 11, pp. 1370 - 1381, Nov. 2010
- [132] R. J. Dengler, et al. "Toward real-time penetrating imaging radar at 670 GHz", *IEEE MTT-S International Microwave Symposium Digest*, pp. 941 - 944, Jun. 2009.

- [133] A. Luukanen, et al. "Developments towards real-time active and passive submillimetre-wave imaging for security applications". *IEEE MTT-S International Microwave Symposium Digest*, pp. 1 - 3, Jun. 2012.
- [134] F. Friederich, et al. "THz active imaging systems with real-time capabilities", *IEEE Trans. THz Sci. Technol.*, vol. 1, no. 1, pp. 183 - 200, Aug. 2011.
- [135] K.B. Cooper. "Imaging, Doppler, and spectroscopic radars from 95 to 700 GHz", *Passive and Active Millimeter-Wave Imaging XIX*, vol. 9830, pp. 983005, May 2016.
- [136] D. A. Robertson, et al. "340 GHz 3D radar imaging test bed with 10 Hz frame rate", *Passive and Active Millimeter-Wave Imaging XV*, vol. 8362, Apr. 2012.
- [137] B. Cheng, et al. "Real-time imaging with a 140 GHz synthetic aperture radar". *IEEE Trans. THz Sci. Technol.*, vol. 3, no. 5, pp. 594 - 605, Sept. 2013.
- [138] D. A. Robertson, T. F. Gallacher, R. Sondena, & D. G. Macfarlane. "Millimeter wave imaging at up to 40 frames per second using an optoelectronic photo-injected Fresnel zone plate lens antenna". *Passive and Active Millimeter-Wave Imaging XIX*, vol. 9830, pp. 983004, May 2016.
- [139] B. Gonzalez-Valdes, et al. "Millimeter wave imaging architecture for on-the-move whole body imaging". *IEEE Trans. Antennas Propag.*, vol. 64, no. 6, pp. 2328 - 2338, Mar. 2016.
- [140] V. Iglesias, J. Grajal, M. A. Sánchez, & M. López-Vallejo, "Implementation of a real-time spectrum analyzer on FPGA platforms". *IEEE Trans. Instrum. Meas.*, vol. 64, no. 2, pp. 338 - 355, Feb. 2015.
- [141] A. Kharin, S. Vityazev, V. Vityazev, & N. Dahnoun. "Parallel FFT implementation on TMS320c66x multicore DSP". *2014 6th European Embedded Design in Education and Research Conference (EDERC)*, pp. 46 - 49, Sept. 2014.
- [142] M. Otten, W. Vlothuizen, H. Spreeuw, & A. Varbanescu. "Real-time processing of multi-channel SAR data with GPUs". *2016 European Radar Conference*, pp. 65 - 68, Oct. 2016.
- [143] NVIDIA - *CUDA C Programming Guide* [Online] <https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html>, January 2017.
- [144] H. Nguyen. *Gpu gems 3*, Addison-Wesley Professional, 2007.
- [145] J. Sanders, & E. Kandrot. *CUDA by example: an introduction to general-purpose GPU programming*, Addison-Wesley Professional, 2010.
- [146] N. Wilt. *The cuda handbook: A comprehensive guide to gpu programming*, Pearson Education, 2013.
- [147] J. W. Cooley, & J. W. Tukey. "An algorithm for machine calculation of complex Fourier series". *Math. Comp.*, vol. 19, pp. 297 - 301, 1965.
- [148] NVIDIA - *Profiler's User Guide* [Online], http://docs.nvidia.com/cuda/pdf/CUDA_Profiler_Users_Guide.pdf
- [149] J. Lobeiras, M. Amor, & R. Doallo. "Designing efficient index-digit algorithms for CUDA GPU architectures". *IEEE Trans. Parallel Distrib. Syst.*, vol. 27, no. 5, pp. 1331 - 1343, May 2016.
- [150] F. Wang, et al. "Time-domain algorithm for FMCW based short distance ranging system". *Antennas and Propagation (EuCAP), 2015 9th European Conference on*, pp. 1-5, May 2015.

- [151] Kitware - *VTK User's Guide* [Online], <https://vtk.org/vtk-users-guide/>
- [152] Pthreads-win32 - POSIX Threads for Win32 [Online], <http://sourceware.org/pthreads-win32/>
- [153] M. Qi, T.-T. Cao, & T.-S. Tan. "Computing 2D constrained Delaunay triangulation using the GPU". *IEEE Trans. Vis. Comput. Graphics*, vol. 19, no. 5, pp. 736 - 748, Mar. 2013
- [154] E. Sintorn, & U. Assarsson. "Fast parallel GPU-sorting using a hybrid algorithm". *Journal of Parallel and Distributed Computing*, vol. 68, no. 10, pp. 1381-1388, Oct. 2008.
- [155] J. Laviada, A. Arboleya-Arboleya, & F. Las-Heras. "Multistatic millimeter-wave imaging by multiview portable camera". *IEEE Access*, vol. 5, pp. 19259-19268, Jul. 2017.
- [156] F. Gumbmann., & S. S. Ahmed. "Walk through screening with multistatic mmW technology" *Millimetre Wave and Terahertz Sensors and Technology IX*, vol. 9993, p. 999306, Oct. 2016.
- [157] X. Yang, et al. "Joint multistatic THz imaging radars for standoff personnel screening". *Journal of Applied Remote Sensing*, vol. 12, no. 3, pp. 034002, Jul. 2018.
- [158] N. Llombart, & B. Blázquez. "Refocusing a THz imaging radar: Implementation and measurements". *IEEE Trans. Antennas Propag.*, vol. 62, no. 3, pp. 1529-1534, Mar. 2014.
- [159] D. A. Robertson, et al. "High resolution, wide field of view, real time 340GHz 3D imaging radar for security screening". *Passive and Active Millimeter-Wave Imaging XX*, vol. 10189, May 2017.
- [160] M. A. Patrick, et al. "Range resolved mode mixing in a large volume for the mitigation of speckle and strategic target orientation requirements in active millimeter-wave imaging". *JOSA A*, vol. 32, no. 4, pp. 637-646, Apr. 2015.
- [161] M. A. Patrick, J. A. Holt, C. D. Joye, & F. C. De Lucia. "Elimination of speckle and target orientation requirements in millimeter-wave active imaging by modulated multimode mixing illumination". *JOSA A*, vol. 29, no. 12, pp. 2643-2656, Dec. 2012.
- [162] C. am Weg, et al. "Fast active THz-camera with global illumination". *Infrared, Millimeter, and Terahertz Waves, 2009. IRMMW-THz 2009. 34th International Conference on*, pp. 1-2, Sept. 2009.
- [163] W. von Spiegel, C. am Weg, R. Henneberger, R. Zimmermann, H. G. & Roskos. "Illumination aspects in active terahertz imaging". *IEEE Trans. Microw. Theory Tech.*, vol. 58, no. 7, pp. 2008-2013, Jul. 2010.
- [164] B. González Valdés, et al. "Physical optics simulation of a THz standoff imaging system". *2018 IEEE International Symposium on Antennas and Propagation and USNC-URSI Radio Science Meeting APS/URSI 2018*, July 2018.
- [165] L. Pérez-Eijo, et al. "On the use of a hybrid monostatic-bistatic radar system for a 3D standoff imaging", *12th European Conference on Antennas and Propagation (EuCAP 2018)*, Apr. 2018.
- [166] G. Zhao, et al. "Terahertz dielectric properties of polystyrene foam." *JOSA B*, vol. 19, no. 6, pp. 1476-1479, Jun. 2002.
- [167] C. R. Dietlein, et al. "Absorption, transmission, and scattering of expanded polystyrene at millimeter-wave and terahertz frequencies". *Passive Millimeter-Wave Imaging Technology XI*, vol. 6948, Apr. 2008.

- [168] S. S. Ahmed, S. Sayed, A. Schiessl, & L.-P. Schmidt. "Illumination properties of multistatic planar arrays in near-field imaging applications". *The 7th European Radar Conference*, pp. 29-32, Sept. 2010.
- [169] R. A. Newcombe, et al. "KinectFusion: Real-time dense surface mapping and tracking". *2011 IEEE International Symposium on Mixed and Augmented Reality*, pp. 127-136, Oct. 2011.
- [170] P. J. Besl, & N. D. McKay. "Method for registration of 3-D shapes". *Sensor Fusion IV: Control Paradigms and Data Structures*, vol. 1611, pp. 586-607, Apr. 1992.
- [171] J. Barber, et al. "Developing an ANSI standard for image quality tools for the testing of active millimeter wave imaging systems". *Passive and Active Millimeter-Wave Imaging XX*, vol. 10189, May 2017.
- [172] Z. Wang, & A. C. Bovik. "Modern image quality assessment". *Synthesis Lectures on Image, Video, and Multimedia Processing*, vol. 2, no. 1, pp. 1-156. Dec. 2006.
- [173] Y. Cocheril, & R. Vauzelle. "A new ray-tracing based wave propagation model including rough surfaces scattering". *Progress In Electromagnetics Research*, vol. 75, pp. 357-381, 2007.
- [174] L. Z. Stec, F. J.W. Podd, & A. J. Peyton. "A methodology for the optimisation of a mm-wave scanner". *Millimetre Wave and Terahertz Sensors and Technology IX*, vol. 9993, Oct. 2016.
- [175] T. Möller, & B. Trumbore. "Fast, minimum storage ray/triangle intersection". *ACM SIGGRAPH 2005 Courses*, Jul. 2005
- [176] F. T. Ulaby, R. K. Moore, & A. K. Fung. *Microwave remote sensing: Active and passive. Vol. II. Radar remote sensing and surface scattering and emission theory*, Addison-Wesley, 1982.
- [177] K. Tomiyasu. "Relationship between and measurement of differential scattering coefficient and bidirectional reflectance distribution function (BRDF)". *IEEE Trans. Geosci. Remote Sens.*, vol. 26, no. 5, pp. 660-665, Sept. 1988.
- [178] J.F. Blinn, "Models of light reflection for computer synthesized pictures". *ACM SIGGRAPH computer graphics*, vol. 11, no. 2, pp. 192-198, Jul. 1977.
- [179] R. L. Cook, & K. E. Torrance. "A reflectance model for computer graphics". *ACM Transactions on Graphics (TOG)*, vol. 1, no. 1, pp. 7-24 Jan. 1982.
- [180] D. M. Sheen, A. M. Jones, & T. E. Hall. "Simulation of active cylindrical and planar millimeter-wave imaging systems". *Passive and Active Millimeter-Wave Imaging XXI*, vol. 10634, May 2018.
- [181] R. Appleby, & S. Ferguson. "Submillimetre wave imaging and security: imaging performance and prediction". *Millimetre Wave and Terahertz Sensors and Technology IX*, vol. 9993, Oct. 2016.
- [182] J. Laviada, et al. "Multiview three-dimensional reconstruction by millimetre-wave portable camera". *Scientific reports*, vol. 7, no. 1, pp. 6479, Jul. 2017.
- [183] T.-T. Cao, et al. "A GPU accelerated algorithm for 3D Delaunay triangulation". *Proceedings of the 18th meeting of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pp. 47-54, Mar. 2014.
- [184] A. Nanjappa. *Delaunay Triangulation in R3 on the GPU*. NUS, 2012.

- [203] G. M. Brooker. "Understanding millimetre wave FMCW radars". *1st International Conference on Sensing Technology*, pp. 152-157, Nov. 2005.
- [204] M. Pauli, et al. "Miniaturized millimeter-wave radar sensor for high-accuracy applications". *IEEE Trans. Microw. Theory Techn.*, vol. 65, no. 5, pp. 1707 - 1715, Mar. 2017.
- [205] M. A. Richards, J. Scheer, W. A. Holm, & W. L. Melvin. *Principles of modern radar*, SciTech Pub., 2010.
- [206] M. C. Budge, & M. P. Burt. "Range correlation effects in radars". *Radar Conference, 1993., Record of the 1993 IEEE National*, pp. 212-216, 1993.
- [207] A. B. Constantine. *Antenna theory: analysis and design*, John Wiley & Sons, 2005.
- [208] R. E. Fischer, et al. *Optical system design*, McGraw Hill, 2000.
- [209] J. Luomahaara, J. Säily, P. Pursula, A. Luukanen, & J. Hassel. "Characterization of sub-range-bin spatial features with frequency modulated continuous wave radar operated at 650 GHz", *2016 9th Global Symposium on Millimeter-Waves (GSMM)*, Jul. 2016.
- [210] S. Scherr, et al. "Miniaturized 122 GHz ISM band FMCW radar with micrometer accuracy". *Radar Conference (EuRAD), 2015 European*, pp. 277-280, Sept. 2015.
- [211] S. Scherr, et al. "Parametric estimation of distances with a 122 GHz FMCW radar system". *Microwaves for Intelligent Mobility (ICMIM), 2016 IEEE MTT-S International Conference on*, pp. 1-4, May 2016.
- [212] S. Dill, M. Peichl, E. Schreiber, & H. Anglberger. "Improved characterization of scenes with a combination of MMW radar and radiometer information". *Passive and Active Millimeter-Wave Imaging XX*, vol. 10189, p. 1018907, May 2017.
- [213] H. B. Wallace, & M. J. Rosker. "Analytical performance comparison of active and passive SMMW imaging for contraband detection". *Millimetre Wave and Terahertz Sensors and Technology II*, vol. 7485, p. 74850E, Sept. 2009.
- [214] I. N. Dolganova, et al. "Combined terahertz imaging system for enhanced imaging quality". *Optical and Quantum Electronics*, vol. 48, no. 6, pp. 325, May 2016.
- [215] Brijot Millimeter Wave Technologies Corp. - Product Details [Online], <http://www.brijot.com/?l=products&id=2>
- [216] J. F. Lampin, et al. "MILLIPRISM: millimeter wave passive radiometric imaging system". *8ème Workshop Interdisciplinaire sur la Sécurité Globale*, 2014.
- [217] J. N. Gollub, et al. "Large metasurface aperture for millimeter wave computational imaging at the human-scale". *Scientific Reports*, no. 7, 42650, pp. 1-9, Feb. 2017.
- [218] J. Laviada, et al. "Multiview mm-wave imaging with augmented depth camera information". *IEEE Access*, vol. 6, pp. 16869-16877, Mar. 2018.
- [219] WAVECamm Corp. - Product Details [Online], <http://wavecamm.com/product.php>
- [220] H. Sarbolandi, D. Lefloch, & A. Kolb. "Kinect range sensing: Structured-light versus time-of-flight Kinect". *Computer Vision and Image Understanding*, vol. 139, pp. 1 - 20, Oct. 2015.
- [221] Farran Technology, Ltd. - Imaging Front End [Online], <http://www.farran.com/contentFiles/productFiles/PMMW-10-0001.pdf>

- [222] Flann Microwave Ltd. - Standard Gain Horns Series 240 [Online], <http://flann.com/wp-content/uploads/2016/01/Series-240.pdf>
- [223] A. Bleiweiss, & M. Werman. "Fusing time-of-flight depth and color for real-time segmentation and tracking". *Dynamic 3D imaging*, pp. 58 - 69, Sept. 2009
- [224] R. Appleby. "Passive millimetre-wave imaging and how it differs from terahertz imaging". *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* vol. 362, no. 1815, pp. 379 - 393, Feb. 2004.
- [225] M. Kowalski, & M. Kastek. "Comparative studies of passive imaging in terahertz and mid-wavelength infrared ranges for object detection". *IEEE Trans. Inf. Forensics Security*, vol. 11, no. 9, pp. 2028 - 2035, May 2016.
- [226] J. Shotton, et al. "Real-time human pose recognition in parts from single depth images". *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pp. 116 - 124, Aug. 2011.
- [227] H. Li, B. S. Manjunath, & S. K. Mitra. "Multisensor image fusion using the wavelet transform". *Graphical Models and Image Processing*, vol. 57, no. 3, pp. 235 - 245, May 1995.
- [228] S. Mallat. *A wavelet tour of signal processing*, Academic press, 1999.
- [229] I. Daubechies. *Ten lectures on wavelets*. vol. 61. SIAM, 1992.
- [230] G. Pajares, & J. M. De La Cruz. "A wavelet-based image fusion tutorial". *Pattern recognition*, vol. 37, no. 9, pp. 1855 - 1872, Sept. 2004.
- [231] Silicon Radar GmbH - Products Under Development [Online], <https://siliconradar.com/products/products-under-development/>
- [232] PDWT - Parallel DWT [Online], <https://github.com/pierrepaleo/PDWT>
- [233] H.-M. Chen, et al. "Imaging for concealed weapon detection: a tutorial overview of development in imaging sensors and processing". *IEEE Signal Process. Mag.*, vol. 22, no. 2, pp. 52-61, Mar. 2005.
- [234] A. Frimpong. "Introduction of full body image scanners at the airports: a delicate balance of protecting privacy and ensuring national security". *Journal of Transportation Security*, vol. 4, no. 3, pp. 221-229, Sept. 2011.
- [235] J. F. Tugus. "Privacy and body scanners at EU airports". *International Federation for Information Processing*, vol. 49, pp. 49-54, 2013.
- [236] M. Leese. "'We were taken by surprise': body scanners, technology adjustment, and the eradication of failure". *Critical Studies on Security*, vol. 3, no. 3, pp. 269-282, Sept. 2015.
- [237] J. Deng, et al. "ImageNet: A large-scale hierarchical image database". *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248-255, Jun. 2009.
- [238] Passenger Screening Algorithm Challenge - Kaggle [Online], <https://www.kaggle.com/c/passenger-screening-algorithm-challenge>
- [239] S. Yeom, et al. "Real-time concealed-object detection and recognition with passive millimeter wave imaging". *Optics express*, vol. 20, no. 9, pp. 9371-9381, Apr. 2012.

- [240] S. López-Tapia, R. Molina, & N. Pérez de la Blanca. "Using machine learning to detect and localize concealed objects in passive millimeter-wave images". *Engineering Applications of Artificial Intelligence*, vol. 67, pp. 81-90, Jan. 2018.
- [241] S. López Tapia, R. Molina, & N. Pérez de la Blanca. "Detection and localization of objects in Passive Millimeter Wave Images". *2016 24th European Signal Processing Conference (EUSIPCO)*, pp. 2101-2105, Aug. 2016.
- [242] O. Martínez, et al. "Concealed object detection and segmentation over millimetric waves images". *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops*, pp. 31-37, Jun. 2010.
- [243] E. Gonzalez-Sosa, et al. "Exploring body shape from mmw images for person recognition". *IEEE Trans. Inf. Forensics Security*, vol. 12, no. 9, pp. 2078-2089, Sept. 2017.
- [244] S. Madhogaria, & M. Schikora. "Automatic detection of hidden threats in the TeraSCREEN passive millimeter-wave imaging subsystem". *Passive and Active Millimeter-Wave Imaging XVIII*, vol. 9462, pp. 94620J, May 2015.
- [245] H. Isiker, et al. "An autotclassification procedure for concealed weapon detection in millimeter-wave radiometric imaging systems". *Microwave and Optical Technology Letters*, vol. 60, no. 3, pp. 583-594, Mar. 2018.
- [246] I. Gómez Maqueda, et al. "Fast millimeter wave threat detection algorithm". *2015 23rd European Signal Processing Conference (EUSIPCO)*, pp. 599-603, Aug. 2015.
- [247] D.-S. Lee, et al. "Image segmentation of concealed objects detected by passive millimeter wave imaging". *2009 34th International Conference on Infrared, Millimeter, and Terahertz Waves*, pp. 1-2, Sept. 2009.
- [248] X. Shen, et al. "Detection and segmentation of concealed objects in terahertz images". *IEEE Trans. Image Process.*, vol. 17, no. 12, pp. 2465-2475, Dec. 2008.
- [249] S. López-Tapia, R. Molina, & N. Pérez de la Blanca. "Deep CNNs for object detection using passive millimeter sensors". *IEEE Trans. Circuits Syst. Video Technol.*, Nov. 2017.
- [250] B. Kumar, P. Sharma, & D. Singh. "Development of scale and rotation invariant neural network based technique for detection of dielectric contrast concealed targets with millimeter wave system". *Defence Science Journal*, vol. 67, no. 6, pp. 674-681, Nov. 2017.
- [251] L. Carrer, & A. G. Yarovoy. "Concealed weapon detection using UWB 3-D radar imaging and automatic target recognition". *The 8th European Conference on Antennas and Propagation (EuCAP 2014)*, pp. 2786-2790, Apr. 2014.
- [252] P. E. Keller, et al. "Use of a neural network to identify man-made structure in millimeter-wave images for security screening applications". *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, pp. 2009-2014, Jul. 2006.
- [253] M. Grega, et al. "Automated detection of firearms and knives in a CCTV image". *Sensors*, vol. 16, no. 1, pp. 47, Jan. 2016.
- [254] M. Parande, & S. Soma. "Concealed weapon detection in a human body by infrared imaging", *International Journal of Science and Research*, vol. 4, no. 9, pp. 182-188, Sept. 2015.
- [255] S. Yuenyong, N. Hnoohom, & K. Wongpatikaseree. "Automatic detection of knives in infrared images". *2018 International ECTI Northern Section Conference on Electrical, Electronics, Computer and Telecommunications Engineering (ECTI-NCON)*, pp. 65-68, Feb. 2018.

- [256] S. MacIntosh, et al. "Detection of person borne IEDs using multiple cooperative sensors". *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense X*, vol. 8019, pp. 801910, Jun. 2011.
- [257] S. Nercessian, K. Panetta, & S. Agaian. "Automatic detection of potential threat objects in X-ray luggage scan images". *2008 IEEE Conference on Technologies for Homeland Security*, pp. 504-509, May 2008.
- [258] D. H. Hubel, & T. N. Wiesel. "Receptive fields of single neurones in the cat's striate cortex". *The Journal of physiology*, vol. 148, no. 3, pp. 574-591, Oct. 1959.
- [259] F. Rosenblatt. "The perceptron: a probabilistic model for information storage and organization in the brain". *Psychological review*, vol. 65, no. 6, pp. 386, Nov. 1958.
- [260] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, & L. D. Jackel. "Backpropagation applied to handwritten zip code recognition". *Neural computation*, vol. 1, no. 4, pp. 541-551, Dec. 1989.
- [261] Y. LeCun, L. Bottou, Y. Bengio, P. & Haffner. "Gradient-based learning applied to document recognition". *Proc. IEEE*, vol. 86, no. 11, pp. 2278-2324, Nov. 1998.
- [262] Y. LeCun, Y. Bengio, & G. Hinton. "Deep learning". *Nature*, vol. 521, no. 7553, pp. 436, May 2015.
- [263] A. Krizhevsky, I. Sutskever, & G. E. Hinton. "ImageNet classification with deep convolutional neural networks". *Advances in neural information processing systems*, pp. 1097-1105, 2012.
- [264] V. Nair, & G. E. Hinton. "Rectified linear units improve restricted boltzmann machines." *Proceedings of the 27th international conference on machine learning (ICML-10)*, pp. 807-814, 2010.
- [265] I. Goodfellow, Y. Bengio, & A. Courville. *Deep learning*. MIT press, 2016.
- [266] A. Sharif Razavian, et al. "CNN features off-the-shelf: an astounding baseline for recognition". *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 806-813, 2014.
- [267] Y. Yao, L. Rosasco, & A. Caponnetto. "On early stopping in gradient descent learning". *Constructive Approximation*, vol. 26, no. 2, pp. 289-315, Aug. 2007.
- [268] H. He, & E. A. Garcia. "Learning from imbalanced data". *IEEE Trans. Knowl. Data Eng.*, vol. 9, pp. 1263-1284, Dec. 2008.
- [269] L. Perez, & J. Wang. "The effectiveness of data augmentation in image classification using deep learning". *arXiv preprint arXiv:1712.04621*, Dec. 2017.
- [270] A. Bhattacharyya, C. H. Lind, & R. Shirpurkar. *Threat Detection in TSA Scans using AlexNet*, University of California, San Diego, 2018.
- [271] A. A. R. Guimaraes, & G. Tofighi. "Detecting zones and threat on 3D body in security airports using deep learning machine". *arXiv preprint arXiv:1802.00565*, Feb. 2018.
- [272] A. de Jorge Sánchez. *Study of Machine Learning Algorithms To Detect Threats In Airport Passengers*, Universitat Politècnica de Catalunya, Bachelor's Thesis, 2018.
- [273] K. Simonyan, & A. Zisserman. "Very deep convolutional networks for large-scale image recognition". *arXiv preprint arXiv:1409.1556*, Sept. 2014.

- [274] FLIR Systems - E60 Infrared Camera [Online], <https://www.flir-direct.com/product/flir-e60-infrared-camera>
- [275] A. Ellmauthaler, C. L. Pagliari, E. A. da Silva, J. N. Gois, & S. R. Neves. "A visible-light and infrared video database for performance evaluation of video/image fusion methods", *Multidimensional Systems and Signal Processing*, vol. 30, no. 1, pp. 119-143, 2019.
- [276] Object Tracking and Classification Beyond the Visible Spectrum (OTCBVS) - OTCBVS Benchmark Dataset Collection [Online], <http://vcip1-okstate.org/pbvs/bench/>
- [277] M. Kristoffersen, J. Dueholm, R. Gade, & T. Moeslund. "Pedestrian counting with occlusion handling using stereo thermal cameras", *Sensors*, vol. 16, no. 1, pp. 62, Jan. 2016.
- [278] A. Berg, J. Ahlberg, & M. Felsberg. "A thermal object tracking benchmark", *12th IEEE International Conference on Advanced Video and Signal Based Surveillance*, pp. 1-6, Aug. 2015.
- [279] M. Felsberg, et al. "The thermal infrared visual object tracking VOT-TIR2015 challenge results", *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 76-88, 2015.
- [280] Z. Wu, N. Fuller, D. Theriault, & M. Betke. "A thermal infrared video benchmark for visual analysis". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 201-208, 2014.
- [281] W. Treible, et al. "Cats: A color and thermal stereo benchmark", *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2961-2969, 2017.
- [282] E. Gebhardt, & M. Wolf. "Camel dataset for visual and thermal infrared multiple object detection and tracking", *15th IEEE International Conference on Advanced Video and Signal Based Surveillance*, pp. 1-6, Nov. 2018.
- [283] Q. Liu, X. Lu, Z. He, C. Zhang, & W. S. Chen. "Deep convolutional neural networks for thermal infrared object tracking", *Knowledge-Based Systems*, vol. 134, pp. 189-198, 2017.
- [284] O. Russakovsky, J. Deng, Z. Huang, A. C. Berg, & L. Fei-Fei. "Detecting avocados to zucchinis: what have we done, and where are we going?", *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2064-2071, 2013.
- [285] M. Andriluka, et al. "2D human pose estimation: New benchmark and state of the art analysis". *Proceedings of the IEEE Conference on computer Vision and Pattern Recognition*, pp. 3686-3693, 2014.
- [286] K. Soomro, A. R. Z., & M. Shah. "UCF101: A dataset of 101 human actions classes from videos in the wild". *arXiv preprint arXiv:1212.0402*, Dec. 2012.
- [287] M. J. Ackerman. "The visible human project". *Proc. IEEE*, vol. 86, no. 3, pp. 504-511, Mar. 1998.
- [288] C. Szegedy, et al. "Going deeper with convolutions". *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1-9, 2015.
- [289] K. He, et al. "Deep residual learning for image recognition". *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770-778, 2016.

- [290] J. Hu, L. Shen, & G. Sun. "Squeeze-and-excitation networks". *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132-7141, 2018.
- [291] Keras - Keras Documentation [Online], <https://keras.io/#why-this-name-keras>
- [292] TensorFlow - TensorFlow Guide [Online], <https://www.tensorflow.org/guide>
- [293] Caffe - Caffe [Online], <http://caffe.berkeleyvision.org/>
- [294] NVIDIA - DIGITS [Online], <https://developer.nvidia.com/digits>
- [295] H. Su, et al. "Multi-view convolutional neural networks for 3D shape recognition". *Proceedings of the IEEE international conference on computer vision*, pp. 945-953, 2015.
- [296] K. S. C. Jaikla. *Threats detection for airport body scan images*, Stanford University, 2018.
- [297] Kaggle - Post Competition Architecture Discussion [Online], <https://www.kaggle.com/c/passenger-screening-algorithm-challenge/discussion/45805>
- [298] M. Kastek, M. Kowalski, H. Polakowski, P. Lagueux, & M. A. Gagnon. "Passive signatures concealed objects recorded by multispectral and hyperspectral systems in visible, infrared and terahertz range", *Active and Passive Signatures V*, vol. 9082, p. 90820A, Jun. 2014.
- [299] C. U. Grosse, & M. Ohtsu. *Acoustic emission testing*. Springer Science & Business Media, 2008.
- [300] J. P. Hou, et al. "Prediction of impact damage in composite plates". *Composites Science and Technology*, vol. 60, no. 2, pp. 273-281, Feb. 2000.
- [301] R. Zoughi. *Microwave non-destructive testing and evaluation principles. Vol. 4*. Springer Science & Business Media, 2012.
- [302] S. Kharkovsky, & R. Zoughi. "Microwave and millimeter wave nondestructive testing and evaluation - Overview and recent advances". *IEEE Instrum. Meas. Mag.*, vol. 10, no. 2, pp. 26-38 Apr. 2007.
- [303] C. Viegas, et al. "Active millimeter-wave radiometry for Nondestructive Testing/Evaluation of composites-glass fiber reinforced polymer". *IEEE Trans. Microw. Theory Tech.*, vol. 65, no. 2, pp. 641-650, Feb. 2017.
- [304] S. Becker, T. Ullmann, & G. Busse. "3D Terahertz imaging of hidden defects in oxide fibre reinforced ceramic composites". *Proc. 4th International Symposium on NDT in Aerospace*. 2012.
- [305] M. Nezhad, J. Adametz, & L.-P. Schmidt. "Wideband imaging systems in the mm-Wave and THz range for security and nondestructive testing". *General Assembly and Scientific Symposium (URSI GASS), 2014 XXXIth URSI*, pp. 1-4, Aug. 2014.
- [306] D. Bourreau, A. Péden, & S. Le Maguer. "A quasi-optical free-space measurement setup without time-domain gating for material characterization in the W-band". *IEEE Trans. Instrum. Meas.*, vol. 55, no. 6, pp. 2022-2028, Dec. 2006.
- [307] A. Kazemipour, et al. "Design and calibration of a compact quasi-optical system for material characterization in millimeter/submillimeter wave domain". *IEEE Trans. Instrum. Meas.*, vol. 64, no. 6, pp. 1438-1445, Jun. 2015.

- [308] B. M. Fischer, et al. "Investigating material characteristics and morphology of polymers using terahertz technologies". *IEEE Trans. THz Sci. Technol.*, vol. 3, no. 3, pp. 259-268, Jun. 2015.
- [309] N. Palka, et al. "Non-destructive testing of polyethylene composite by terahertz radiation". *Microwave, Radar and Wireless Communications (MIKON), 2016 21st International Conference on*, pp. 1-4, May 2016.
- [310] A. Badolato Martín, "System for the automatic guidance of vehicles by means of dielectric changes in a pre-recorded guide-rail", *IPC: G05D1/02 (2006.01)*, Nov. 2016.
- [311] A. Pinker, & C. Smith. "Vulnerability of the GPS Signal to Jamming". *GPS Solutions*, vol. 3, no. 2, pp. 19-27, Oct. 1999.
- [312] A. B. Hillel, et al. "Recent progress in road and lane detection: a survey". *Machine vision and applications*, vol. 25, no. 3, pp. 727-745, Apr. 2014.
- [313] M. S. Iftekhar, N. Saha, & Y. M. Jang. "Stereo-vision-based cooperative-vehicle positioning using OCC and neural networks". *Optics Communications*, vol. 352, pp. 166-180, Oct. 2015.
- [314] G. Yadav, S. Maheshwari, & A. Agarwal. "Fog removal techniques from images: A comparative review and future directions". *Signal Propagation and Computer Technology (ICSPCT), 2014 International Conference on*, pp. 44-52, Jul. 2014.
- [315] J. Levinson, M. Montemerlo, & S. Thrun. "Map-based precision vehicle localization in urban environments". *Robotics: Science and Systems*, vol. 4, Jun. 2007.
- [316] T. G. Phillips, N. Guenther, & P. R. McAree. "When the dust settles: The four behaviors of LIDAR in the presence of fine airborne particulates". *Journal of Field Robotics*, vol. 34, no. 5, pp. 985-1009, Aug. 2017.
- [317] G. Kim, J. Eom, & Y. Park. "Investigation on the occurrence of mutual interference between pulsed terrestrial LIDAR scanners". *Intelligent Vehicles Symposium (IV), 2015 IEEE*, pp. 437-442, Jun. 2015.
- [318] M. Nezaal, J. Schür, & L.-P. Schmidt. "Non-destructive testing of glass fibre reinforced plastics with a synthetic aperture radar in the lower THz region". *Infrared, Millimeter, and Terahertz Waves (IRMMW-THz), 2012 37th International Conference on*, pp. 1-2, Sept. 2012.
- [319] M. Nezaal, J. Schür, & L.-P. Schmidt. "Non-destructive testing of glass fibre reinforced plastics with a full polarimetric imaging system". *Infrared, Millimeter, and Terahertz waves (IRMMW-THz), 2014 39th International Conference on*, pp. 1-2, Sept. 2014.
- [320] J. C. Weatherall, et al. "The identification of explosives in millimeter-wave imaging systems". *Passive and Active Millimeter-Wave Imaging XXI*, vol. 10634, May 2018.
- [321] J. Laviada, et al. "Nondestructive evaluation of microwave-penetrable pipes by synthetic aperture imaging enhanced by full-wave field propagation model". *IEEE Trans. Instrum. Meas.*, vol. 99, pp. 1-8, Aug. 2018.
- [322] R. Zoughi, et al. "Millimeter wave imaging as a tool for traceability and identification of tattooed markers in leather". *2018 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, pp. 1-5, May 2018.
- [323] J. R. Gallion, & R. Zoughi. "Millimeter-wave imaging of surface-breaking cracks in steel with severe surface corrosion". *IEEE Trans. Instrum. Meas.*, vol. 66, no. 10, pp. 2789-2791, Oct. 2017.

- [324] I. Mehdipour, et al. "Use of near-field microwave reflectometry to evaluate steel fiber distribution in cement-based mortars". *Journal of Materials in Civil Engineering*, vol. 29, no. 7, Mar. 2017.
- [325] F. Friederich, et al. "Industrial radome inspection with terahertz waves". *Millimetre Waves and Terahertz Technologies (UCMMT), 2017 10th UK-Europe-China Workshop on*, pp. 1-2, Sept. 2017.
- [326] J. P. Guillet, et al. "Art painting diagnostic before restoration with terahertz and millimeter waves". *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 38, no. 4, pp. 369-379, Apr. 2017.
- [327] M. Ghaani, C. A. Cozzolino, G. Castelli, & S. Farris. "An overview of the intelligent packaging technologies in the food sector". *Trends in Food Science & Technology*, vol. 51, pp. 1 - 11, May 2016.
- [328] A. Regattieri, M. Gamberi, & R. Manzini. "Traceability of food products: General framework and experimental evidence". *Journal of Food Engineering*, vol. 81, no. 2, pp. 347 - 356, Jul. 2007.
- [329] G. Roussos. "Enabling RFID in retail". *Computer*, vol. 39, no. 3, pp. 25 - 30, Mar. 2006.
- [330] C. Xin. "RFID applications in retail industry". *Internation Colloquium on Computing, Communication, Control, and Management*, Sept. 2009.
- [331] M. Bustillo. "Wal-Mart radio tags to track clothing", *The Wall Street Journal*, Jul. 2010.
- [332] I. K. Tzoulis, Z. S. Andreopoulou, & E. Voulgaridis. "Wood Tracking Information Systems to Confront Illegal Logging". *Journal of Agricultural Informatics*, vol. 5, no. 1, pp. 9 - 17, Jul. 2014.
- [333] P. A. Wäger, M. Eugster, L. M. Hilty, & C. Som. "Smart labels in municipal solid waste - a case study for the Precautionary Principle?". *Environmental Impact Assessment Review*, vol. 25, no. 5, pp. 567 - 586, Jul. 2005.
- [334] R. M. Hornby. "RFID solutions for the express parcel and airline baggage industry". *IEEE Colloquium on RFID technology*, Oct. 1999.
- [335] A. Furness. "Machine-readable data carriers - a brief introduction to automatic identification and data capture". *Assembly Automation*, vol. 20, no. 1, pp. 28 - 34, Mar. 2000.
- [336] R. Want, "An introduction to RFID technology". *IEEE Pervasive Comput.*, vol. 5, no. 1, pp. 25 - 33, Feb. 2006.
- [337] S. Mori, H. Nishida, & H. Yamada. *Optical Character Recognition*, John Wiley & Sons, Inc., 1999.
- [338] B. Duong, H. Liu, L. Ma, & M. Su. "Covert thermal barcodes based on phase change nanoparticles". *Nature Scientific Reports*, vol. 4, Jun. 2014.
- [339] P. M. Reyes, S. Li, & J. K. Visich. "Determinants of RFID adoption stage and perceived benefits". *European Journal of Operational Research*, vol. 254, no. 3, Nov. 2016.
- [340] N. Huber, K. Michael, & L. McCathie. "Barriers to RFID adoption in the supply chain". *IEEE RFID Eurasia*, pp. 1 - 6, Sept. 2007.

- [341] G. White, G. Gardiner, G. P. Prabhakar, & A. Abd Razak. "A comparison of barcoding and RFID technologies in practice". *Journal of Information, Information Technology and Organizations*, vol. 2, pp. 119 - 132, 2007.
- [342] S. Shin, & B. Eksioglu. "Effects of RFID technology on efficiency and profitability in retail supply chains". *Journal of Applied Business Research*, vol. 30, no. 3, pp. 633 - 646, May 2014.
- [343] S. Shin, & J. T. Tucci. "Wal-Mart's dilemma in the 21st century: sales growth vs. inventory growth". *Journal of Applied Business Research*, vol. 31, no. 1, pp. 37 - 46, Jan. 2015.
- [344] Y. Yang, M. Mandehgar, & D. R. Grischkowsky. "Broadband THz signals propagate through dense fog". *IEEE Photon. Technol. Lett.*, vol. 27, no. 4, pp. 383 - 386, Feb. 2015.
- [345] K. Su, L. Moeller, R. B. Barat, & J. F. Federici, "Experimental comparison of terahertz and infrared data signal attenuation in dust clouds". *Journal of the Optical Society of America A*, vol. 29, no. 11, pp. 2360 - 2366, Nov. 2012.
- [346] M. N. Afsar. "Precision millimeter-wave measurements of complex refractive index, complex dielectric permittivity, and loss tangent of common polymers". *IEEE Trans. Instrum. Meas.*, vol. 36, no. 2, pp. 530 - 536, Jun. 1987.
- [347] A. Juels. "RFID security and privacy: a research survey". *IEEE J. Sel. Areas Commun.*, vol. 24, no. 2, pp. 381 - 394, Feb. 2006
- [348] T. Lotlikar, R. Kankapurkar, A. Parekar, & A. Mohite. "Comparative study of Bar-code, QR-code and RFID System". *International Journal of Computer Technology and Applications*, vol. 4, no. 5, pp. 817-821, Sept. 2013.
- [349] S. K. Timalina, R. Bhusal, & S. Moh, S. "NFC and its application to mobile payment: Overview and comparison". *Information Science and Digital Content Technology (ICIDT), 2012 8th International Conference on*, vol. 1, pp. 203-206.
- [350] M. Topkara, T. Erickson, U. Topkara, & C. Narayanaswami. "Enabling multiple QR codes in close proximity". *arXiv preprint arXiv:1510.08210*, Oct. 2015.
- [351] I. Baker. *Fifty Materials That Make the World*, Springer, 2018.
- [352] D. A. DiGiovanni, et al. "Backscattering of ground terrain and building materials at submillimeter-wave and terahertz frequencies". *Passive and Active Millimeter-Wave Imaging XVI*, vol. 8715, May 2013.
- [353] K. Kaiser. *Transmission Lines, Matching, and Crosstalk*, Taylor and Francis, 2006
- [354] R. A. Maleki, G. Meiser. "Managing returnable containers logistics - A case study part II - Improving visibility through using automatic identification technologies". *International Journal of Engineering Business Management*, vol. 3, no. 2, pp. 45-54, May 2011
- [355] Motorola - Managing Returnable Assets in Motion: RFID for Returnable Transport Items [Online], <http://www.bluestarinc.com/fileadmin/bluestar-typo3/documents/PDF-CaseStudies/Motorola/Moto-Returnable.pdf>
- [356] Y. Khan, et al. "Monitoring of vital signs with flexible and wearable medical devices". *Advanced Materials*, vol. 28, no. 22, pp. 4373-4395, Jun. 2016.
- [357] H. Rajagopalan, & Y. Rahmat-Samii. "Novel ingestible capsule antenna designs for medical monitoring and diagnostics". *Antennas and Propagation (EuCAP), 2010 Proceedings of the Fourth European Conference on*, pp. 1-5, Apr. 2010.

- [358] C. Li, et al. "Radar remote monitoring of vital signs". *IEEE Microw. Mag.*, vol. 10, no. 1, pp. 47-56, Feb. 2009.
- [359] L. Tarassenko, et al. "Non-contact video-based vital sign monitoring using ambient light and auto-regressive models". *Physiological measurement*, vol. 35, no. 5, pp. 807, Mar. 2014.
- [360] A. K. Abbas, et al. "Neonatal non-contact respiratory monitoring based on real-time infrared thermography". *Biomedical engineering online*, vol. 10, no. 1, pp. 93, Dec. 2011.
- [361] M. Baboli, et al. "Good night: Sleep monitoring using a physiological radar monitoring system integrated with a polysomnography system". *IEEE Microw. Mag.*, vol. 16, no. 6, pp. 34-41, Jul. 2015.
- [362] N. Hafner, et al. "Non-contact cardiopulmonary sensing with a baby monitor". *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pp. 2300-2302, Aug. 2007.
- [363] H.-C. Kuo, et al. "A fully integrated 60-GHz CMOS direct-conversion Doppler radar RF sensor with clutter canceller for single-antenna noncontact human vital-signs detection". *IEEE Trans. Microw. Theory Tech.*, vol. 64, no. 4, pp. 1018-1028, Apr. 2016.
- [364] J. Tu, T. Hwang, & Jenshan Lin. "Respiration rate measurement under 1-D body motion using single continuous-wave Doppler radar vital sign detection system". *IEEE Trans. Microw. Theory Tech.*, vol. 64, no. 6, pp. 1937-1946, Jun. 2016.
- [365] V. R. Radzicki, et al. "Standoff CW radar for through-the-wall detection of human heartbeat signatures". *Radar Conference (RadarConf), 2016 IEEE*, pp. 1-6, May 2016.
- [366] M. Kagawa, K. Suzumura, & T. Matsui. "Sleep stage classification by non-contact vital signs indices using Doppler radar sensors". *Engineering in Medicine and Biology Society (EMBC), 2016 IEEE 38th Annual International Conference of the*, pp. 4913-4916, Aug. 2016.
- [367] M. Forouzanfar, et al. "Event recognition for contactless activity monitoring using phase-modulated continuous wave radar". *IEEE Trans. Biomed. Eng.*, vol. 64, no. 2, pp. 479-491, Feb. 2017.
- [368] G. Wang, et al. "A hybrid FMCW-interferometry radar for indoor precise positioning and versatile life activity monitoring". *IEEE Trans. Microw. Theory Techn.*, vol. 62, no. 11, pp. 2812-2822, Nov. 2014.
- [369] T.-Y. Huang, L. F. Hayward, & J. Lin. "Non-invasive measurement and analysis of laboratory rat's cardiorespiratory movement". *IEEE Trans. Microw. Theory Techn.*, vol. 65, no. 2, pp. 574-581, Feb. 2017.
- [370] A. Singh, et al. "Activity monitoring and motion classification of the lizard *Chamaeleo jacksonii* using multiple Doppler radars". *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pp. 4525-4528, Aug. 2012.
- [371] N. Hafner, J. C. Drazen, & V. M. Lubecke. "Fish heart rate monitoring by body-contact Doppler radar". *IEEE Sensors J.*, vol. 13, no. 1, pp. 408-414, Jan. 2013.
- [372] X. Huang, et al. "Real-time non-contact infant respiratory monitoring using UWB radar". *Communication Technology (ICCT), 2015 IEEE 16th International Conference on*, pp. 493-496, Oct. 2015.

- [373] S. Suzuki, et al. "An approach to a non-contact vital sign monitoring using dual-frequency microwave radars for elderly care". *Journal of Biomedical Science and Engineering*, vol. 6, no. 07, pp. 704, Jul. 2013.
- [374] Novelda AS - XeThru Products [Online], <https://shop.xethru.com/xethru-products>
- [375] SpecOps Group Inc. - X3FINDER [Online], <https://www.specopsgroup.com/finder>
- [376] Neteera Technologies Ltd. - Neteera Chip [Online], <https://www.neteera.com/product>
- [377] Texas Instruments Inc. - Driver vital sign detection demonstration using mmWave radar sensors [Online], <https://training.ti.com/driver-vital-sign-detection-demonstration-using-mmwave-radar-sensors#>
- [378] A. Santra, et al. "Short-range multi-mode continuous-wave radar for vital sign measurement and imaging". *Radar Conference (RadarConf18)*, 2018 IEEE, pp. 0946-0950, Apr. 2018.
- [379] Qualcomm Tricorder XPRIZE - Competition [Online], <https://tricorder.xprize.org/prizes/tricorder>
- [380] Children's National - Finalists Selected fro \$250k Pediatric Medical Device Competition [Online], <https://childrensnational.org/news-and-events/childrens-newsroom/2016/finalists-selected-for-250k-pediatric-medical-device-competition>
- [381] Innovation4Health - 2018 Competition [Online], <https://www.innovation4health.com/events-1/>
- [382] G. Wang, et al. "Application of linear-frequency-modulated continuous-wave (LFMCW) radars for tracking of vital signs". *IEEE Trans. Microw. Theory Tech.*, vol. 62, no. 6, pp. 1387-1399, Jun. 2014.
- [383] A. D. Droitcour, O. Boric-Lubecke, & G. T. A. Kovacs. "Signal-to-noise ratio in Doppler radar system for heart and respiratory rate measurements". *IEEE Trans. Microw. Theory Tech.*, vol. 57, no. 10, pp. 2498-2507, Oct. 2009.
- [384] A. Prat, S. Blanch, A. Aguasca, J. Romeu, & A. Broquetas. "Collimated beam FMCW radar for vital sign patient monitoring". *IEEE Trans. Antennas Propag.*, (Early Access), Dec. 2018.
- [385] O. Boric-Lubecke, et al. *Doppler radar physiological sensing*. John Wiley & Sons, 2015.
- [386] L. S. Lilly. *Pathophysiology of heart disease: a collaborative project of medical students and faculty*. Lippincott Williams & Wilkins, 2012.
- [387] A. Gacek, & W. Pedrycz. *ECG signal processing, classification and interpretation: a comprehensive framework of computational intelligence*. Springer Science & Business Media, 2011.
- [388] E. A. Ashley, & J. Niebauer. *Cardiology explained*. Remedica, 2004.
- [389] V. N. Batchvarov, M. Malik, & A. J. Camm. "Incorrect electrode cable connection during electrocardiographic recording". *Europace*, vol. 9, no. 11, pp. 1081-1090, Oct. 2007.
- [390] A. Aladin, et al. "Relation of resting heart rate to risk for all-cause mortality by gender after considering exercise capacity (the Henry Ford exercise testing project)". *The American journal of cardiology*, vol. 114, no. 11, pp. 1701-1706, Dec. 2014.

- [391] Z. Lin, et al. "Heart rate estimation using wrist-acquired photoplethysmography under different types of daily life motion artifact". *Communications (ICC), 2015 IEEE International Conference on*, pp. 489-494, Jun. 2015.
- [392] J. Allen. "Photoplethysmography and its application in clinical physiological measurement". *Physiological measurement*, vol. 28, no. 3, Feb. 2007.
- [393] A. Martínez, et al. "Micrometric deformation imaging at W-Band with GBSAR". *European Journal of Remote Sensing*, vol. 49, no. 1, pp. 719-733, Jan. 2016.
- [394] R. Shriram, M. Sundhararajan, & N. Daimiwal. "Cranial PPG Brain Signal based Cardiovascular Parameter Estimation". *XII Control Instrumentation System Conference (CISCON)*, Nov. 2015.
- [395] R. Piesiewicz, et al. "Properties of building and plastic materials in the THz range". *International Journal of Infrared and Millimeter Waves*, vol. 28, no. 5, pp. 363-371, May 2007.
- [396] Y.-S. Jin, G.-J. Kim, & S.-G. Jeon. "Terahertz dielectric properties of polymers". *Journal of the Korean Physical Society*, vol. 49, no. 2, pp. 513-517, 2006.
- [397] H. Yi, et al. "3-D printed millimeter-wave and terahertz lenses with fixed and frequency scanned beam". *IEEE Trans. Antennas Propag.*, vol. 64, no. 2, pp. 442-449, Feb. 2016.
- [398] R. Sauleau, & B. Barès. "A complete procedure for the design and optimization of arbitrarily shaped integrated lens antennas". *IEEE Trans. Antennas Propag.*, vol. 54, no. 4, pp. 1122-1133, Apr. 2006.
- [399] M. Lauer, & D. Stein. "A train localization algorithm for train protection systems of the future". *IEEE Trans. Intell. Transp. Syst.*, vol. 16, no. 2, pp. 970-979, Apr. 2015.
- [400] Y. Wu, et al. "Vulnerabilities, attacks, and countermeasures in balise-based train control systems". *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 4, pp. 814-823, Apr. 2017.
- [401] S. Dhahbi, et al. "Study of the high-speed trains positioning system: European signaling system ERTMS/ETCS". *Logistics (LOGISTIQUA), 2011 4th International Conference on*, pp. 468-473, May 2011.
- [402] F. Bohringer, & A. Geistler. "Comparison between different fusion approaches for train-borne location systems". *Multisensor Fusion and Integration for Intelligent Systems, 2006 IEEE International Conference on*, pp. 267-272, Sept. 2006.
- [403] D. J. Daniels. *Ground penetrating radar. Vol. 1*. IET, 2004.
- [404] A. A. Vertiy, H. Cetinkaya, & M. Tekbas. "Subsurface sensing at sub-terahertz and terahertz frequencies". *Proceedings of the XIII International Conference on Ground Penetrating Radar*, pp. 1-4, 2010.
- [405] T.-Y. J. Kao, et al. "Design and analysis of a 60-GHz CMOS Doppler micro-radar system-in-package for vital-sign and vibration detection". *IEEE Trans. Microw. Theory Tech.*, vol. 61, no. 4, pp. 1649-1659, Apr. 2013.
- [406] J. Tu, & Jenshan Lin. "Fast acquisition of heart rate in noncontact vital sign radar measurement using time-window-variation technique". *IEEE Trans. Instrum. Meas.*, vol. 65, no. 1, pp. 112-122, Jan. 2016.
- [407] M.-C. Huang, et al. "A self-calibrating radar sensor system for measuring vital signs". *IEEE Trans. Biomed. Circuits Syst.*, vol. 10, no. 2, pp. 352-363, Apr. 2016.

- [408] Z. Peng, et al. "A portable FMCW interferometry radar with programmable low-IF architecture for localization, ISAR imaging, and vital sign tracking". *IEEE Trans. Microw. Theory Tech.* vol. 65, no. 4, pp. 1334-1344, Apr. 2017.
- [409] L. Ren, et al. "Phase-based methods for heart rate detection using UWB impulse Doppler radar". *IEEE Trans. Microw. Theory Tech.*, vol. 64, no. 10, pp. 3319-3331, Oct. 2016.
- [410] J. Lin, & W. Wu. "Vital sign radars: Past, present, and future". *Wireless and Microwave Technology Conference (WAMICON), 2014 IEEE 15th Annual*, pp. 1-4, Jun 2014.
- [411] C. Li, & J. Lin. "Random body movement cancellation in Doppler radar vital sign detection". *IEEE Trans. Microw. Theory Tech.*, vol. 56, no. 12, pp. 3143-3152, Dec. 2008.
- [412] C. Gu, et al. "Doppler radar vital sign detection with random body movement cancellation based on adaptive phase compensation". *Microwave Symposium Digest (IMS), 2013 IEEE MTT-S International*, pp. 1-3, Jun. 2013.
- [413] J. M. Muñoz-Ferreras, et al. "Random body movement mitigation for FMCW-radar-based vital-sign monitoring". *Biomedical Wireless Technologies, Networks, and Sensing Systems (BioWireleSS), 2016 IEEE Topical Conference on*, pp. 22-24, Jan. 2016.
- [414] A. Procházka, et al. "Microsoft Kinect visual and depth sensors for breathing and heart rate analysis". *Sensors*, vol. 16, no. 7, pp. 996, Jul. 2016.
- [415] W. Verkrusye, L. O. Svaasand, & J. S. Nelson. "Remote plethysmographic imaging using ambient light". *Optics express*, vol. 16, no. 26, pp. 21434-21445, Dec. 2008.
- [416] C. Li, Y. Xiao, & J. Lin. "Experiment and spectral analysis of a low-power *Ka*-band heartbeat detector measuring from four sides of a human body". *IEEE Trans. Microw. Theory Tech.*, vol. 54, no. 12, pp. 4464-4471, Dec. 2006.
- [417] K. Tian, J. Li, & X. Yang. "A novel method of micro-Doppler parameter extraction for human monitoring terahertz radar network". *Ad Hoc Networks*, vol. 58, pp. 222-230, Apr. 2017.
- [418] C. Suh, & Y.-B. Ko. "Design and implementation of intelligent home control systems based on active sensor networks". *IEEE Trans. Consum. Electron.*, vol. 54, no. 3, pp. 1177-1184, Aug. 2008.
- [419] M. Mubashir, L. Shao, & L. Seed. "A survey on fall detection: Principles and approaches". *Neurocomputing*, vol. 100, pp. 144-152, Jan. 2013.
- [420] F. Chen, et al. "Detection of the vibration signal from human vocal folds using a 94-GHz millimeter-wave radar". *Sensors*, vol. 17, no. 3, pp. 543, 2017.
- [421] R. Hu, & B. Raj. "A robust voice activity detector using an acoustic Doppler radar". *Automatic Speech Recognition and Understanding, 2005 IEEE Workshop on*, pp. 319-324, Nov. 2005.
- [422] M. Pieraccini, et al. "High-speed CW step-frequency coherent radar for dynamic monitoring of civil engineering structures". *Electronics Letters*, vol. 40, no. 14, pp. 907-908, Jul. 2008.
- [423] V. C. Chen, et al. "Micro-Doppler effect in radar: phenomenon, model, and simulation study". *IEEE Transactions on Aerospace and electronic systems*, vol. 42, no. 1, pp. 2-21, Jan. 2006.

- [424] A. Frischen, J. Hasch, & C. Waldschmidt. "FMCW ramp non-linearity effects and measurement technique for cooperative radar". *Radar Conference (EuRAD), 2015 European*, pp. 509-512, Sept. 2015.
- [425] A. Frischen, J. Hasch, & C. Waldschmidt. "Contour recognition with a cooperative distributed radar sensor network". *Microwaves for Intelligent Mobility (ICMIM), 2015 IEEE MTT-S International Conference on*, pp. 1-4, Apr. 2015.
- [426] A. Frischen, et al. "A low-phase-noise 122-GHz FMCW radar sensor for distributed networks". *Radar Conference (EuRAD), 2016 European*, pp. 49-52, Oct. 2016.
- [427] A. Frischen, J. Hasch, & C. Waldschmidt. "A cooperative MIMO radar network using highly integrated FMCW radar sensors". *IEEE Trans. Microw. Theory Tech.*, vol. 65, no. 4, pp. 1355-1366, Apr. 2017.
- [428] K. Lorincz, et al. "Sensor networks for emergency response: challenges and opportunities". *IEEE Pervasive Comput.*, vol. 4, pp. 16-23, Oct. 2004.
- [429] E. Onur, et al. "Surveillance wireless sensor networks: Deployment quality analysis". *IEEE Netw.*, vol. 21, no. 6, Nov. 2007.
- [430] I. Colomina, & P. Molina. "Unmanned aerial systems for photogrammetry and remote sensing: A review". *ISPRS Journal of photogrammetry and remote sensing*, vol. 92, pp. 79-97, Jun. 2014.
- [431] A. Okabe, B. Boots, K. Sugihara. *Spatial tessellations: Concepts and applications of Voronoi diagrams*, Wiley, 1992
- [432] M. Ramírez Torres. *Diseño y caracterización de la etapa de generación de señal para un radar de imagen de 3D*. Master's Thesis, ETSIT-UPM, 2016.
- [433] R. Xu, & D. Wunsch. "Survey of clustering algorithms". *IEEE Trans. Neural Netw.*, vol. 16, no. 3, pp. 645-678, May 2005.
- [434] M. Ester, et al. "A density-based algorithm for discovering clusters in large spatial databases with noise". *KDD*, vol. 96, no. 34. Aug. 1996.
- [435] H. Zhou, P. Wang, & H. Li. "Research on adaptive parameters determination in DBSCAN algorithm". *Journal of Information & Computational Science*, vol. 9, no. 7, pp. 1967-1973, Jul. 2012.
- [436] D. Arlia, & M. Coppola. "Experiments in parallel clustering with DBSCAN". *European Conference on Parallel Processing*, pp. 326-332, Aug. 2001.
- [437] B. Wang, et al. "Design and optimization of DBSCAN Algorithm based on CUDA". *arXiv preprint arXiv:1506.02226*, Jun. 2015.
- [438] B. Welton, E. Samanas, & B. P. Miller. "Mr. Scan: extreme scale density-based clustering using a tree-based network of GPGPU nodes". *High Performance Computing, Networking, Storage and Analysis (SC), 2013 International Conference for*, pp. 1-11, Nov. 2013.
- [439] B. Welton, & B. P. Miller. "The anatomy of Mr. Scan: a dissection of performance of an extreme scale GPU-based clustering algorithm". *Proceedings of the 5th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems*, pp. 54-60, Nov. 2014.
- [440] M. Ankerst, et al. "OPTICS: ordering points to identify the clustering structure". *ACM Sigmod record*, vol. 28. no. 2, pp. 49-60, Jun. 1999.

- [441] C. T. Zahn, & R. Z. Roskies. "Fourier descriptors for plane closed curves". *IEEE Trans. Comput.*, vol. 100, no. 3, pp. 269-281, Mar. 1972.
- [442] H. Bay, T. Tuytelaars, & L. Van Gool. "Surf: Speeded up robust features". *European conference on computer vision*, pp. 404-417, May 2006.
- [443] D. G. Lowe, "Object recognition from local scale-invariant features". *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, vol. 99, no.2, pp. 1150-1157, Sept. 1999.
- [444] J. B.T.M. Roerdink, & Arnold Meijster. "The watershed transform: Definitions, algorithms and parallelization strategies". *Fundamenta informaticae*, vol. 41, no. 1, pp. 187-228, Jan. 2000.
- [445] J. Canny. "A computational approach to edge detection". *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, no. 6, pp. 679-698, Nov. 1986.
- [446] R. Schnabel, R. Wahl, & R. Klein. "Efficient RANSAC for point-cloud shape detection". *Computer graphics forum*, vol. 26, no. 2, pp. 1-12, May 2007.
- [447] J. Tremblay, et al. "Training deep networks with synthetic data: Bridging the reality gap by domain randomization". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops 2018* pp. 969-977, 2018.
- [448] S. Karanam, et al. "A systematic evaluation and benchmark for person re-identification: Features, metrics, and datasets". *arXiv preprint arXiv:1605.09653*, May 2016.
- [449] Z. Wu, et al. "A thermal infrared video benchmark for visual analysis". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 201-208, 2014.
- [450] P. Jagalingam, & A. V. Hegde. "A review of quality metrics for fused image". *Aquatic Procedia*, vol. 4, pp. 133-142, Jan. 2015.
- [451] S. S. Ahmed, A. Schiessl, & L.-P. Schmidt. "A novel fully electronic active real-time imager based on a planar multistatic sparse array". *IEEE Trans. Microw. Theory Tech.*, vol. 59, no. 12, pp. 3567-3576, Dec. 2011.
- [452] Z. Wang, et al. "Image quality assessment: from error visibility to structural similarity". *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600-612, Apr. 2004.
- [453] Silicon Radar GmbH - EvalKits [Online], http://www.siliconradar.com/evalkits_e.html
- [454] S. M. Al-Salem, P. Lettieri, & J. Baeyens. "Recycling and recovery routes of plastic solid waste (PSW): A review". *Waste management*, vol. 29, no. 10, pp. 2625-2643, Oct. 2009.
- [455] R. A. Fenner, E. J. Rothwell, & L. L. Frasc. "A comprehensive analysis of free-space and guided-wave techniques for extracting the permeability and permittivity of materials using reflection-only measurements". *Radio Science*, vol. 47, no. 1, pp. 1-13, Feb. 2012.
- [456] I. Dunayevskiy, et al. "Millimeter-and submillimeter-wave characterization of various fabrics". *Applied optics*, vol. 46, no. 24, pp. 6161-6165, Aug. 2007.
- [457] K. B. Cooper, J. F. Trabert, & R. J. Dengler. "Carrier noise-limited penetration in THz radar imaging". *Microwave Symposium Digest (MTT), 2012 IEEE MTT-S International*, pp. 1-3, Jun. 2012.

-
- [458] D. M. Sheen, et al. "Active imaging at 350 GHz for security applications". *Passive Millimeter-Wave Imaging Technology XI*, vol. 6948, Apr. 2008.
- [459] D. Nuessler. "THz imaging for recycling of black plastics". *Microwave Conference (GeMIC), 2014 German*, pp. 1-4, Mar. 2014.
- [460] D. M. Hailu, & D. Saeedkia. "Novel applications of terahertz technology for the plastic industry". *Infrared, Millimeter, and Terahertz waves (IRMMW-THz), 2016 41st International Conference on*, pp. 1-2, Sept. 2016.
- [461] S.-B. Roh, S.-K. Oh, & W. Pedrycz. "Identification of black plastics based on fuzzy rbf neural networks: Focused on data preprocessing techniques through fourier transform infrared radiation". *IEEE Trans. Ind. Informat.*, vol. 14, no. 5, pp. 1802-1813, May 2018.
- [462] Fraunhofer - blackValue [Online], <https://www.blackvalue.de/en.html>
- [463] U. C. Hasar, et al. "Two-step numerical procedure for complex permittivity retrieval of dielectric materials from reflection measurements". *Applied Physics A*, vol. 116, no. 4, pp. 1701-1710, Sept. 2014.
- [464] J. M. Catalá-Civera, et al. "A simple procedure to determine the complex permittivity of materials without ambiguity from reflection measurements". *Microwave and Optical Technology Letters*, vol. 25, no. 3, pp. 191-194, May 2000.

Annexes

Annex A

Supplementary Data for a 300 GHz Imaging Radar

A.1 Image Representation Considerations

The scan pattern shown in Fig. 2.3 results in a 3-D point cloud of (x, y, z) detections, which the radar can then display in several representations. On the processing side of the radar, every spot of the point cloud is associated to a single received chirp signal, and its spectrum is calculated via an FFT. To be detected and represented, each potential spot (a single chirp) in the scan pattern must surpass a double threshold boundary, where the first threshold is an absolute power level (i.e. a hard threshold based on a constant dBm minimum to limit noise and reduce false alarms) and the second is a variable dynamic range (i.e. a soft threshold based on a limit dependent on the maximum power detected and the side-lobe levels associated with that peak). The power of the spectral peak associated to each 3-D coordinate must exceed the maximum of the two thresholds mentioned. An illustration of the double threshold technique applied to a received spectrum is shown in Fig. A.1.

Figure A.1: Double threshold power boundary applied to a sample chirp spectrum.

Of course, the magnitude of these thresholds affects the final imaging results. Selecting a noise threshold too low will allow false alarms to appear in the 3-D representation as noise

around the target, but rising it too high will create empty areas in the FoV where detections are not powerful enough. As for the dynamic range, selecting a lower value will permit the maximum peak's side lobes to be detected, creating duplicate range layers in the image, but setting it too high will forbid all peaks less powerful than the maximum from being represented. An illustration of these effects is shown in Fig. A.2, where the same scan of a female mannequin hiding a handgun beneath a T-shirt is presented three times, but each time the thresholds have been set to different values.

Figure A.2: Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a female mannequin is hiding a handgun underneath a T-shirt, when the noise and dynamic range thresholds are (a) -89 dBm and 12 dB, (b) -79 dBm and 12 dB, and (c) -79 dBm and 30 dB, respectively. Colorbar: Range in meters.

In addition to spatial location, each pixel in the point cloud can contain abundant data regarding the FoV spot it represents, such as its associated beat frequency, its range to the radar, its returned power, or its phase. To enhance visual information further than only 3-D coordinates, texture (through a synthetic colormap) can be applied to the displayed point cloud. To do this, a colormap is associated to one of the mentioned spot variables (e.g. range or returned power), and each spot is colored according to its magnitude. For example, in Fig. A.3 the scenario from Fig. A.2 is repeated, but in each representation the colormap is associated to different spot information (range, returned power, and phase).

Traditionally, scans from security imaging radars have been presented in either range or returned power colormaps. Although it was initially more common in the literature to display results in returned power representations, these are now being progressively switched to range-only images. This is due to the large dependence of received power on incident angle when the wavelength is very small (as in the 300 GHz band), where smooth surfaces that produce specular reflection can even cause targets to vanish from representation (e.g. up to 70 dB losses have been encountered with the 300 GHz imager when tilting a metallic plate just 3° [97]). This means that returned power distributions from the FoV have high variance from one measurement to another, and thus shapes can be heavily distorted, causing false threat detections or misses. As illustration of this effect, a metallic plate held with a tripod at the standoff range is imaged in a sequence where 1° rotations in the y -axis are performed, with the set-up shown in Fig. A.4. The results can be seen in Fig. A.5, which proves that even small shifts in orientation can cause threats to deform or even vanish.

In addition to the choice of texture, further display options are available for the captured data. Regarding the selected spectrum peaks to be displayed, two decisions must be made:

Figure A.3: Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a female mannequin is hiding a handgun underneath a T-shirt when the colormap is associated to (a) range (in m), (b) returned power (in dBm), and (c) phase (in degrees).

Figure A.4: Imaging scenario for the orientation sequence measured in Fig. A.5, where a smooth metallic plate is placed 8 m away from the radar using a tripod.

Figure A.5: Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a smooth metallic plate is imaged with (a) no rotation, (b) 1° rotation, and (c) 2° rotation around the y -axis. Colorbar: Received power in dBm.

the maximum number of peaks to be selected, and the order in which these peaks are sorted. Because each datapoint in the spectrum corresponds to a different distance, selecting multiple peaks from a single chirp spectrum is equivalent to adding more range layers to the represented image. This is useful when the target is wearing various clothing layers, to represent all of them simultaneously in a single image. An illustration of the effect of selecting additional layers on the image, for the scenario in Fig. A.2, is shown in Fig. A.6.

Figure A.6: Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a female mannequin is hiding a handgun underneath a T-shirt when (a) one peak, (b) two peaks, and (c) three peaks are selected. Colorbar: Range in meters.

Meanwhile, the order in which the peaks are selected influences the final resulting image. Three hierarchies for peak sorting are considered for the 300 GHz imager: first-to-last in range, last-to-first in range, and first-to-last in returned power. Because the number of peaks to be represented is limited, applying each of these hierarchies results in different image information. When peaks are ordered first-to-last in range the display contains the outer

range layers of the FoV (i.e. clothing on the target), while the inner layers (i.e. skin and hidden threats) are displayed when sorting last-to-first in range. Meanwhile, sorting the peaks in a first-to-last fashion regarding returned power places importance on the range layers with the most powerful reflections, such as metallic objects. A comparison of the resulting images when applying this criterion and selecting only one spectrum peak, for the imaging scenario in Fig. A.2, is shown in Fig. A.7.

Figure A.7: Resulting 3-D point clouds from the 300 GHz imaging radar of a scenario where a female mannequin is hiding a handgun underneath a T-shirt when the (a) first range peak, (b) last range peak, and (c) maximum power peak are chosen. Colorbar: Range in meters.

As can be expected, selecting last-to-first in range exposes the handgun more clearly, because it is located farther away than the T-shirt. If the most powerful peaks are selected, it can be seen that in some spots these correspond to the T-shirt rather than the target's body, but on the other hand the gun is now denser (because it is a metallic object that produces strong reflections).

For the 300 GHz imaging radar, several displays are also available along with the point cloud representations here presented. A collection of all the displays implemented is shown in Fig. A.8 and Fig. A.9, for the case of range-based texture, in the imaging scenario of Fig. A.2. In Fig. A.8(b) a rectangular grid covering the entire FoV in 1 mm steps is generated, and the point cloud information is interpolated into the grid to fill out missing areas through synthetic data. Since the interpolating function can have abrupt variations, a 2-D symmetric median filter can be applied to the interpolated grid [such as in Fig. A.8(c)] to smooth out the data. This filter serves an additional purpose: the mm-level step size of the grid is much smaller than the cm-level spatial resolution of the radar, so the filter's kernel size is adjusted to the imager's resolution in order to smooth the grid into shapes comparable to the radar's real results. In Fig. A.9(a), instead of generating interpolated data, the point cloud coordinates are triangulated in 3-D to form the facets of a polygonal surface. This is performed through the implementation of a Delaunay triangulation [431], and the motive is to approximately reconstruct the actual surface of the target. For better visualization, the facets of the triangles created are colored with their corresponding texture in Fig. A.9(c).

Within the possible surface reconstructions, several processing options are also available. In the 2-D media filtering, the kernel size is a parameter that can be modified as desired. If a small kernel size is chosen, the filtering will preserve the interpolation's features, at the cost of producing a rough surface not alike what the radar would actually generate. On the other hand, a very large kernel size will create a smoother layer, but will also deform the image by enveloping too much area in each filtering step. An illustration of these considerations is presented in Fig. A.10, where the imaging results from the scenario in Fig. A.2 are filtered with varying kernel size. As for the Delaunay-based reconstruction, a similar effect can

Figure A.8: Collection of representations available in the 300 GHz imaging radar for the results of scanning a female mannequin hiding a handgun underneath a T-shirt. (a) is a point cloud display of the 3-D coordinates, (b) is the result of applying a cubic interpolation of the point cloud into a rectangular grid occupying the FoV with a 1 mm step size, and (c) is the output from 2-D filtering the previous interpolated grid with a median kernel to smooth its surface. Colorbar: Range in meters.

Figure A.9: Collection of representations available in the 300 GHz imaging radar for the results of scanning a female mannequin hiding a handgun underneath a T-shirt. (a) is a point cloud display of the 3-D coordinates, (b) is the result of applying a Delaunay triangulation in 3-D to the point cloud, where each triangle is colored according to its range, and (c) is that same reconstruction but also providing texture to the facets of the triangles generated. Colorbar: Range in meters.

Figure A.10: Examples of interpolated and 2-D filtered outputs for a scan at 300 GHz of a female mannequin hiding a handgun underneath a T-shirt, where the filter's kernel sizes are (a) $1 \times 1 \text{ cm}^2$, (b) $2.5 \times 2.5 \text{ cm}^2$, and (c) $4 \times 4 \text{ cm}^2$. Colorbar: Range in meters.

be accomplished with a division of the coordinates by a compression factor in any axis, performed before the triangulation, and which must be undone when representing the results. Because during the Delaunay triangulation coordinates near to each other are clustered into triangles, compressing an axis is equivalent to facilitating the generation of triangles in that axis's direction. A sample sequence is given in Fig. A.11, where the imaging results from the scenario in Fig. A.2 are compressed in one axis or another to produce varying reconstruction results. As can be noticed, compressing in only one axis produces a surface with clear edges in that direction, so the compression has to be balanced in all dimensions to generate a realistic surface.

Figure A.11: Examples of reconstructed surfaces for a scan at 300 GHz of a female mannequin hiding a handgun underneath a T-shirt, using a 3-D Delaunay triangulation, where the previous to the triangulation the point cloud's axes are compressed by factors of (a) 1.3 in x only, (b) 1.3 in x and 1.6 in y , and (c) 1.6 in y only. Colorbar: Range in meters.

A.2 Automatic Target Recognition and Clustering

As mentioned in Chapter 4, the current trend in security screening systems is to ease the workload imposed on human operators by implementing ATD algorithms. With the appropriate processing chains, hidden threats can be automatically identified within an image without the need of a human's visual confirmation. Aside from the increased performance in throughput and detection rate, these implementations can have added benefits, such as less violation of privacy.

To reach this goal, deep learning approaches like the one implemented in Chapter 4 can be adapted to the 300 GHz radar's images. For this implementation a re-training of the neural network with images from the 300 GHz system would be necessary, and some minor adjustments to the network architecture could be useful to emphasize the different image features at the higher frequency. In addition to the adaptation of the neural network, several pre-processing stages can be developed to ease threat detection.

Within these stages, clustering techniques [433] (which can group together 3-D coordinates into clusters that share common properties) have proven to be visually effective in highlighting concealed objects in the 300 GHz radar's point cloud images. Clustering can be useful to accentuate the threat in an image, thereby facilitating the network's detection or even classification. Ideally, the previous clustering stage would group all pixels of the threat into a single cluster, while keeping all other pixels (body, clothing, non-threatening items) in one or more different clusters.

One of the more common clustering techniques, the density-based spatial clustering of applications with noise (DBSCAN) algorithm [434], performs exceptionally well in point clouds deteriorated with noise (such as those from the 300 GHz imaging radar). The DBSCAN technique consists of separating useful points from noise based on the density of the potential cluster: if the evaluated point is near to a large number of other points, they are all labeled as a single cluster, whereas if the point is alone or near few other point they are considered outlier noise. To do this separation, the algorithm needs several user-determined thresholds to establish the clusters, such as the maximum distance between points or the minimum

number of points. Although the ideal values of these parameters are previously unknown and can vary from one imaging scenario to another, heuristic methods to determine useful settings have been designed [435].

In Fig. A.12, the result of applying the DBSCAN technique on three different imaging scenarios from the 300 GHz radar is presented. It can be noted, from observing the abdominal area in the three scenarios, that the threat is perfectly separated from the subject's body into a different cluster in every case. Although it is true that most of the subjects' bodies are grouped into a single cluster, some body parts (such as faces or breasts) are incorrectly separated. This effect could prove to be a problem, since it could lead the neural network to inaccurately learn to detect these clusters instead of the threats. Nonetheless, the total number of clusters and their grouped coordinates can be modified through variations of the DBSCAN's parameters. This allows reducing erroneous clusters, at the expense of having to individually process each image to set the correct values.

Figure A.12: Resulting 3-D point clouds from the 300 GHz imaging radar, after being processed with the DBSCAN clustering technique, of scenarios where (a) a male human wearing a shirt hides a plastic bottle (colored orange), (b) a female mannequin wearing a T-shirt hides a rectangular simulated explosive device (colored orange), and (c) that same mannequin hiding a gun (colored brown). The colors represent the different clusters detected.

In addition, the DBSCAN algorithm presents numerous parallel-computing variations [436–439]. These are useful because the algorithm needs to loop through all points in the cloud and calculate their distance to all other points, which is a very resource-intensive task but can be computed in parallel for each point. Through these parallel implementations images can be processed with the DBSCAN algorithm at a faster rate. A custom version of the DBSCAN method on CUDA, to be implemented in the GPU platforms described in Section 2.2, has been developed. This code produces equivalent results to the serial version, as seen in Fig. A.13, at less than one hundredth of the processing time. Future work regarding clustering techniques should be focused on performing parametric sweeps to extract the optimal clustering thresholds to accurately extract threats only, as well as compare the DBSCAN method against other clustering approaches (such as the OPTICS algorithm [440], which reduces the need to set specific thresholds by sorting all possible clusters).

Instead of a neural network approach to ATD, other machine learning techniques could also be evaluated. The processing stages for three different strategies for threat detection can be seen in Fig. A.14. In Fig. A.15 the outputs at the different stages of the first approach are represented. In this method, the captured point cloud is first registered with a reference

Figure A.13: Comparison of the clustering results of (a) serial and (b) parallel versions of the DBSCAN algorithm for an imaging scenario of the 300 GHz radar where a female mannequin conceals a gun (colored brown) under a T-shirt.

point cloud (in this case, of the mannequin without any clothing or threats) via the ICP algorithm. After that, the points in the captured cloud that have no correspondence in the reference cloud are extracted. Because the extraction takes into account the range distance, only coordinates protruding from the mannequin's body are extracted. These are coordinates that are only present in the point cloud where the mannequin is hiding a gun, so they are interpreted as potential objects of interest. Finally, the boundary of the extracted coordinates is computed, to confirm whether the corresponding surface resembles a threat.

Figure A.14: Block diagrams of the processing stages for three different ATD strategies for the 300 GHz imaging radar.

The second strategy in Fig. A.14, with the outputs shown in Fig. A.16, shows a similar approach. In this case, instead of using a reference point cloud to extract potential threats, the captured cloud is clustered with the DBSCAN technique. Afterwards, the boundary of the detected clusters is computed, and their Fourier descriptors [441] are compared against

Figure A.15: Outputs of the processing stages for the first threat detection strategy in Fig. A.14. In (a), corresponding to step A in Fig. A.14, the point cloud captured with the 300 GHz imaging radar for a scenario of a female mannequin wearing a T-shirt and hiding a gun is shown. In (b), corresponding to step B in Fig. A.14, the captured point cloud is registered to a reference point cloud of the mannequin with no clothing or threat. In (c), corresponding to step C in Fig. A.14, points belonging to both clouds or not protruding from the body are eliminated, leaving only those from the gun. In (d), corresponding to step D in Fig. A.14, the boundary of the extracted coordinates is computed, showing visual similarity to a gun. Colorbar: Range in meters.

that of a reference shape. For this approach, although no reference point cloud is needed, a large number of reference shapes would be required to compare the clusters against all possible threat types, orientations, sizes, and other variations. To reduce this requirement several feature extractors robust against image transformations, like rotations and scalings, could be used (such as the SURF [442] or SIFT [443] algorithms).

Figure A.16: Outputs of the processing stages for the second threat detection strategy in Fig. A.14. In (a), corresponding to step A in Fig. A.14, the point cloud captured with the 300 GHz imaging radar for a scenario of a female mannequin wearing a T-shirt and hiding a gun is shown. In (b), corresponding to step E in Fig. A.14, the captured point cloud is clustered with the DBSCAN technique. In (c), corresponding to step F in Fig. A.14, the boundary of the gun cluster is computed and visually compared against a reference threat shape. In (e), corresponding to step E in Fig. A.14, the Fourier descriptors of the boundaries are compared in a spectral representation. Colorbar: Range in meters.

The final strategy in Fig. A.14, with the outputs shown in Fig. A.17, is based on a very different approach. For this method, the interpolated and filtered grid representation of the capture data is used. This output is then processed with a watershed algorithm [444] which segments the image according to the different range values, thereby extracting objects that protrude from the mannequin's body. The different segments are then interpreted as

potential threats.

Figure A.17: Outputs of the processing stages for the third threat detection strategy in Fig. A.14. In (a), corresponding to step H in Fig. A.14, the interpolated and filtered grid from the 300 GHz imaging radar for a scenario of a female mannequin wearing a T-shirt and hiding a gun is shown. In (b), corresponding to step I in Fig. A.14, the grid is processed with a watershed algorithm to segment surfaces in different ranges (the segmented surfaces are highlighted with different colors). In (c), corresponding to step J in Fig. A.14, the gun segment is extracted. Colorbar: Range in meters.

In future work, ATD strategies using neural networks should be compared, in terms of detection rate and processing time, against other machine learning techniques such as those presented in this section. In addition to those already described, several other strategies for object detection can be designed, such as stages based on edge detection in an image [445] or shape extractions in point clouds [446].

A.3 Image Database and Quality Metric

To enable the automatic detection algorithms of Section A.2, a large database of images captured by the radar is necessary to adjust their parameters and test their performance. Although this task would have previously been very burdensome due to the high number of required images, combined with the scanning and processing limitations, thanks to the development of the real-time architecture in Section 2.2 images can now be produced as frames from a continuous video. Additionally, the software described in Section 2.3.2 allows generating simulated images that mimic the radar's output without the need to capture new scenarios. Fast generation of images for the database can be enabled with the use of synthetic images, created without the need for capturing real data, as machine learning algorithms are starting to be tested with scenarios fully-rendered by software [447].

In addition to its internal use as training samples for the detection algorithms, the development of this database has tremendous novelty in the security imaging field if made publicly-available. Although some examples of databases from security screening systems exist (such as the radar-based dataset in Chapter 4, the collected videos in [448], or thermal datasets [449]), to our knowledge, the proposed database would be the first public collection of images generated with a 300 GHz active standoff radar.

For this objective, a collection of more than 100 different imaging scenarios has already been put together. The subjects scanned range from plastic mannequins to male and female humans, concealing various threats (such as guns, knives, or explosives) under common clothing layers. In addition, different radar orientations are considered, to offer images from all of the subject's sides. A sample assortment of some of the scenarios available in the collection is shown in Fig. A.18. To build a complete database in the future, this collection would need to be incremented to account for a wider variety of clothing items, as well as different body types and ages.

Figure A.18: Sample radar images from the initial database for the 300 GHz imaging radar. The top row is of a female mannequin (a) wearing no clothing and hiding no threat, (b) wearing a T-shirt and hiding a gun, and (c) wearing a T-shirt and hiding an explosive device. The bottom row is of a male human, covering the same clothing and threat scenarios, in the same order. All images are of the inner range layer. Colorbar: Range in meters.

A related necessity is the development of an objective metric to establish a given image's quality. This function would be useful to determine if an image is of high enough quality to enter the database, but also to compare the radar's results under different scenarios (such as changes in architecture, noise, or bandwidth). All of the radar's images published up to date have been evaluated from a qualitative point of view, determining their quality by judging how the subject's physical features are represented and how the concealed threat is detected.

To evaluate and compare images with a consistent and fair method, a quantitative measure is required.

Although throughout the literature some work can be found on the development of image quality tools specifically for active [171] or fused [450] imagers, most systems still use radar parameters to define the quality available. To this end, figures of merit such as cross-range and range resolutions, SNR, FoV area, or probability of detection are still the most common metrics. With the use of these parameters several experiments can be performed to determine image quality, but these deal with system-level aspects which may not be always extrapolated to every individual image.

In traditional imaging devices (such as cameras) several metrics can be used to determine an output's quality [452]. Nonetheless, in these devices most metrics depend on comparing a noisy or distorted image with its reference picture. In active radar imagers these metrics could be useful to compare an image's relative quality against another image of the same static scenario, but not as a global parameter for all images. To extract a universal indicator of quality, the so-called no-reference (when no undistorted reference image is available) or reduced-reference (when only partial data describing an ideal image are available) methods would have to be used [172].

A.4 Minor Changes to the Original System

Outside of the work and objective for this thesis, but concurrent with its development, several modifications have been made to the original 300 GHz described in its first complete-system publication [102]. These modifications are detailed briefly in this section, to serve as an update on the current state of the imager:

- **Compact TRX.** As explained in a recent publication [123], the radar's RF architecture has been updated to a more compact configuration. In this scheme, the imager simultaneously transmits and receives radiation through only one mm-wave chain. This can be done by taking advantage of the SHM's finite isolation¹ between its RF and LO ports as shown in Fig. A.19, which means that part of the LO signal used for downconversion is leaked to the antenna.

With this architecture, the radar still transmits a signal with 27 GHz of bandwidth centered at 300 GHz, and is fully capable of generating images equivalent to the original bistatic set-up, even though the transmitted signal is roughly 20 dB less powerful, as seen in Fig. A.20. Notice the focused nature of the imager, as the radar reaches its peak received power only at a specific range². An imaging comparison between both architectures is presented in Figs. A.21.

The main benefits of the compact configuration are reduced size, cost and power consumption, but received signal linearity is also improved since the transmitted and LO chirps now travel through the same mm-wave components. In addition, because only one mm-wave chain remains in the front end, the beamsplitter used for isolation is no longer necessary.

- **New x6 multiplier.** Although when mounted in compact configuration the TX mm-wave chain is eliminated, the radar can still function in its bistatic architecture as well. With this in mind, an update on the x6 active multiplier that upconverts the chirp from 15.9 GHz - 17.4 GHz to 95.4 GHz - 104.4 GHz has been performed. The component has been replaced with an upgraded version from the same manufacturer, and is thus

¹With approximately 26 dB of attenuation.

²Registered as farther than 8 m due to the propagation path within the reflector system.

Figure A.19: Block diagram comparison between (a) the original mm-wave architecture (with items to be removed in red) and (b) the compact configuration.

Figure A.20: Received power for a metallic target as its range grows progressively, for the bistatic and compact architectures.

Figure A.21: (a) Female mannequin with a T-shirt hiding a set of fake cylindrical explosives in the stomach area. (a) Inner layer image with the compact architecture. (b) Inner layer image with the original architecture. Colorbar: Range in meters..

still a COTS device. The change is motivated by the multiplier's placement in the mm-wave chain; as the first device in the front end, any power ripples in its bandwidth are propagated and amplified through the following components. In addition, the updated multiplier offers an increment of more than 5 dB in output power for the same input, ultimately producing a more powerful transmitted signal, or even allowing the removal of the MPA³. From an RF point-of-view the comparison between the two versions of the multiplier may be seen in Annex B, while an imaging comparison is shown in Fig. A.22. In this last comparison, the images show similar power distributions but higher received power for the new multiplier, which would allow detection of previously-masked less powerful reflections.

- **New DDS.** The DDS used for signal generation in the imager has also been updated, from a DSM303B model from Euvis Inc. to a AD9914 model from Analog Devices [432], partly motivated by the improved performance of the latter device with a much lower cost (from $\sim 1700\text{€}$ to $\sim 600\text{€}$). While using the same oscillator reference the updated DDS produces power (by 2 dB) and lower variations at its 500 MHz - 1000 MHz output. In addition, spur levels in the side bands are lowered by 10 dB in the AD9914 DDS, which ultimately results in less-powerful spurious frequencies transmitted and received by the radar. As an added advantage the new DDS can also generate the swept bandwidth through its image bands at 1.4 GHz - 1.9 GHz and 3.8 GHz - 4.3 GHz, which could allow elimination of the noisy 4.8 GHz LO signal used to upconvert the DDS's output to 5.3 GHz - 5.8 GHz.
- **New custom power supply.** In an effort to further compact the imager, a custom power supply unit is used to feed all active components of the RF architecture. In its newer version (although larger in size), shown in Fig. A.23, the supply allows for re-configuration of any of its twelve voltage outputs and real-time monitoring of their consumption through a text-based interface. In addition, alarm interruptions and shut-downs can be programmed in case any output exceeds a selected threshold, to protect equipment from short-circuits. To allow all these features, the unit is governed by a

³To benefit from less cost, consumption, and isolation problems.

Figure A.22: Imaging comparison of the maximum-power layer of a nude female mannequin. (a) is the result of using the original x6 multiplier, and (b) is the result of using the updated version. Colorbar: Received power in dBm.

microcontroller board connected to the radar's PC through a USB interface.

Figure A.23: Photographs of the (a) previous and (b) new power supply units for the imaging radar.

Annex B

RF Components for a 300 GHz Imaging Radar

B.1 x6 Active Frequency Multiplier

The documents included in this annex are the specifications of the two versions of the x6 active frequency multiplier used in the 300 GHz imaging radar's TX mm-wave chain, as presented in Section A.4. They are the official datasheet from the manufacturer, Radiometer Physics Gmbh. The documents are organized as follows:

- Datasheet from Radiometer Physics Gmbh for AFM6-110 (original x6 multiplier).
- Datasheet from Radiometer Physics Gmbh for AFM6 75-110 +10 (updated x6 multiplier).

In addition to these documents, the two versions of the x6 multiplier have been characterized for output power throughout their bandwidth, while varying the input power and the measurement point in the mm-wave chain. The measurements have been performed by feeding the x6 multiplier with single frequency tones originally generated by the same DDS and microwave upconversion used as source for the imaging radar. Output power is registered with the use of a Erickson PM4 Power Meter from Virginia Diodes Inc. that measures the temperature variance generated by an absorption of power when referenced against a known baseline (effectively working as a bolometer). The power meter includes a straight waveguide transition to W-band, but must be fitted with a taper for the measurements centered at 300 GHz.

The output power of the x6 multipliers when fed with their nominal input power (7 dBm) is compared in Fig. B.1. As can be seen from the figure, the power ripple is comparable in both cases (being 0.47 dB for the original multiplier and 0.65 dB for the updated version), but the new device generates at least 5.12 dB more of output power for the same input power. This can be beneficial for later stages in the mm-wave chain, where the low efficiencies of the following multiplications drastically reduce the final transmitted power. In addition, a large enough increase in the output power of this first multiplication stage could allow the elimination of the MPA used for W-band amplification (thereby reducing cost, consumption, and adaptation problems). A comparison between the multipliers' responses when they are input with varying power can be seen in Fig. B.2, where both devices show similar decrease in performance when input power is lowered.

Figure B.1: Active x6 multipliers' output power for the imaging radar's 95.4 GHz - 104.4 GHz bandwidth.

Figure B.2: Comparison of output power at the x6 multipliers' output with varying input power (P_{in}). (a) is the original multiplier, and (b) is the updated version.

Signal power is next measured at the TX chain's isolator output, between the x6 multiplier and MPA, in Fig. B.3. In this step, the ripple from the updated device has risen to 0.87 dB, but the original component now also suffers from a 1.22 dB ripple. The lower increment in variation for the updated multiplier can be due to an improved excitation of the following stages in the mm-wave chain because of the increased power, thus offering higher adaptation.

Figure B.3: Isolator's output power for the imaging radar's 95.4 GHz - 104.4 GHz bandwidth.

The next measurement takes place after the MPA, in Fig. B.4. The ripple has increased to 2.70 dB and 3.58 dB for the updated and original multipliers, respectively. In addition to a lesser maximum variation, the new multiplier presents a smoother response with lower local variance. This is once again likely due to the reduction of the standing wave present between the isolator and MPA in the original architecture. In Fig. B.5 signal power at this stage is compared among the two multipliers when varying input power.

Figure B.4: MPA's output power for the imaging radar's 95.4 GHz - 104.4 GHz bandwidth.

The last measurement takes place at the x3 multipliers output (centered at 300 GHz) in Fig. B.6, where the increase in the final transmitted power and flatness are evident. The new multiplier presents a maximum power variation of 13.46 dB, while that figure is 21.36 dB for

the original device. In addition, the maximum output power of the x3 multiplication stage has increased from 5.23 dBm to 7.39 dBm. Fig. B.7 presents the comparison of x3 output power for the two x6 multiplier alternatives when their input power is varied.

Figure B.5: Comparison of power at the MPA's output with varying input power (P_{in}) for the active x6 multipliers. (a) is the original multiplier, and (b) is the updated version.

Figure B.6: MPA's output power for the imaging radar's 95.4 GHz - 104.4 GHz bandwidth.

Figure B.7: Comparison of power at the x3's output with varying input power (P_{in}) for the active x6 multipliers. (a) is the original multiplier, and (b) is the updated version.

Active Frequency Multiplier
AFM-Series
WR10

Product Data Sheet - April 2010 Iss 1 Rev 1

Product Data Sheet:

Active Frequency Multiplier
Order Code: AFM6-110
Serial - No.: 205

Specifications:

AFM6-110	WR10
Frequency Range	75 - 110 GHz
RF - Interface	WR10 with Flange UG387
Multiplication Factor	6
DC-Requirement	450 mA max. @ +7 VDC
RF - IN	12,5 - 18,33 GHz
RF - Input Power	+7 dBm opt. (+10 dBm max)

Output Power [dBm] @ +7 dBm

AFM6 75-110 +10

Active Frequency Multiplier 75-110 GHz +10

Part-No.: 01100037

Product Description

The **AFM6 75-110 +10** is an active frequency multiplier (x6) with a coaxial input and a WR-10 waveguide output.

Product Features:

- Full Band WR10 (75-110GHz)
- output Power: typ. +10 dBm
- No Mechanical Tuners
- Input Power: +7dBm

Technical Specifications	
Output Frequency [GHz]	75 - 110
RF Port	WR10 (UG387/UM)
Output Power [dBm (typ.)]	+10
Input Frequency [GHz]	12.5 - 18.3
Input Connector	2.92 mm female
Input Power [dBm]	+7
DC Voltage [VDC]	+7 / 650mA (typ.)
Max. Input Power [dBm]	+10
Max. DC Voltage [VDC]	+9
Max. Case Temperature	+45 °C

Typical Performance

B.2 Bandpass Filter

This section is dedicated to the characterization of the band-pass filter used in the 300 GHz imaging radar as an anti-aliasing tool to enable real-time processing, as detailed in Section 2.2.3. The filter is the model ZX75-12-S+ from Mini-Circuits, and its official datasheet is also included in this annex.

To measure the filter's S_{21} response a N9010A EXA signal analyzer from Agilent has been used, along with a N9310A signal generator from the same manufacturer. A sweep from 1 MHz to 20 MHz, with 200 sample points and -30 dBm of input power, is configured and input into the filter. The signal is measured at the filter's output with the analyzer at 180 KHz bandwidth resolution. The result may be seen in Fig. B.8. As can be seen in the figure, the insertion losses in the passband (9 MHz - 15 MHz) are between 1 dB and 2 dB, while the losses in a stopband (for instance, >18 MHz) are greater than 30 dB. These values are within the expected ranges provided by the manufacturer's datasheet.

Figure B.8: (a) Measured S_{21} parameter from Mini-Circuit's bandpass filter. (b) Bandpass zoom of the S_{21} measurement. (c) Superior stopband zoom of the S_{21} measurement.

Coaxial
Bandpass Filter

50Ω 9 to 15 MHz

Maximum Ratings

Operating Temperature	-40°C to 85°C
Storage Temperature	-55°C to 100°C
RF Power Input	0.25 W at 25°C

Permanent damage may occur if any of these limits are exceeded.

Outline Drawing (HY1239)

Outline Dimensions (inch/mm)

A	B	C	D	E	F
1.20	.75	.46	1.18	.04	.17
30.48	19.05	11.68	29.97	1.02	4.32
G	H	J	K	L	M
1.00	.50	.35	.18	.106	1.88
25.40	12.70	8.89	4.57	2.69	47.75

Notes:

* M dimension is 2.05 inch (52.07 mm) for HY1238 Case Style.

Features

- High Rejection
- Sharp Insertion Loss roll off
- Shielded case

Applications

- High rejection applications
- Image rejection
- IF signal processing
- Receivers / Transmitters
- Test Lab

ZX75-12+

MALE FEMALE SMA shown

SMA Connectors	Model	Case
INPUT FEMALE	OUTPUT FEMALE	ZX75-12-S+ HY
MALE	FEMALE	ZX75-12M-S+ HY

+RoHS Compliant

The +Suffix identifies RoHS Compliance. See our website for RoHS Compliance methodologies and qualifications.

Bandpass Filter Electrical Specifications (T_{AMB} = 25°C)

CENTER FREQ. (MHz)	PASSBAND (MHz) (Loss < 3.5dB)	STOPBANDS (MHz)				VSWR (:1)		
		Loss > 20dB		Loss > 40dB		Passband		Stopband
Fc	F1 - F2	F3	F4	F5	F6	Typ.	Max.	Typ
12	9 - 15	7.5	18	7	20 - 3000	1.5	2.2	30

Typical Frequency Response

Functional Schematic

Typical Performance Data at 25°C

Frequency (MHz)	Insertion Loss (dB)	VSWR (:1)
0.3	100.39	1737.18
7.0	53.30	62.05
7.5	36.73	35.46
8.0	14.55	9.74
8.2	5.51	2.37
8.5	3.45	2.36
9.0	1.85	1.22
11.0	1.31	1.40
12.0	1.26	1.37
13.0	1.59	1.77
15.0	2.09	1.75
15.7	4.10	2.81
16.0	6.58	4.66
16.6	15.51	11.77
18.0	37.07	30.49
20.0	60.12	59.91
100.0	94.22	248.17
1000.0	67.12	23.49
3000.0	59.03	6.19

Notes
A. Performance and quality attributes and conditions not expressly stated in this specification document are intended to be excluded and do not form a part of this specification document.
B. Electrical specifications and performance data contained in this specification document are based on Mini-Circuit's applicable established test performance criteria and measurement instructions.
C. The parts covered by this specification document are subject to Mini-Circuit's standard limited warranty and terms and conditions (collectively, "Standard Terms"); Purchasers of this part are entitled to the rights and benefits contained therein. For a full statement of the Standard Terms and the exclusive rights and remedies thereunder, please visit Mini-Circuit's website at www.minicircuits.com/MCLStore/terms.jsp

www.minicircuits.com P.O. Box 350166, Brooklyn, NY 11235-0003 (718) 934-4500 sales@minicircuits.com

REV
M12
ED7
ZX7
UR,
170
Pag

Annex C

Synchronization Electronics for a 300 GHz Imaging Radar

C.1 Scanning Electronics

To accurately represent FoV coordinates in the 300 GHz imaging radar, the scanning mirror's movements must be synchronized with IF signal sampling, since the (x, y) tuple of each chirp is extracted from the mirror's elevation and rotation angles (γ, ϕ) . This is done via individual rotation and elevation motors¹ attached to the mirror, as shown in Fig. C.1, which include encoders and controllers² that convert the motors' movements to pulse sequences.

Figure C.1: (a) CAD diagram and (b) photograph of the scanning mirror, its mechanic structure, and the devices (motors and reduction drive) involved in its movement.

The encoders output two main types of signals: index and quad counts. The index signal is a one-per-rotation pulse useful for synchronization, while the quad counts come in channels with differential signals containing continuous pulse trains as the motors moves (providing step-by-step resolution). The encoders are also connected to a custom-made PCB, shown in Fig. C.2, housing a microcontroller³ and some ICs, designed to handle the encoder outputs.

¹The elevation motor is model 35H4B.2,33-019 from Haydon Kerk Motion Solutions, while the rotation motor is model EC-4 Pole 305013 from Maxon Motor.

²An EPOS2 70/10 rotation motor controller from Maxon Motor, and a PCM4806X elevation motor controller from Haydon Kerk Motion Solutions.

³MOD5270 from NetBurner.

Figure C.2: CAD render of the PCB in charge of motor control and encoder handling.

After properly conditioning these pulses they must be sampled and transferred to the PC. This task is performed by the PCB's microcontroller, using its input/output (I/O) pins and an Ethernet connection. Then, within the polynomial kernel of the CUDA code, each pair of motor angles is attached to the z coordinate from a frequency ramp. Simple diagrams illustrating this entire sequence and the devices involved are shown in Fig. C.3 and Fig. C.4, respectively.

Figure C.3: Block diagram presenting the processes and subsystems involved in motor control and positioning.

Scanning requires a start-up sequence to ensure that the mirror begins the scan at a known position. The details of this sequence are as follows:

- First, the user configures the scan details (such as speed, duration, and sampling frequency) through the software's graphic interface.
- Afterwards, the PC initiates a blank VTK viewing window and the necessary memory allocations for the entire scans. Meanwhile, the microcontroller waits for an Ethernet connection and confirmation message from the PC.
- When the PC finishes the initialization it connects to the microcontroller and orders it to begin the start-up. This sequence is governed by the microcontroller, which communicates data to and from the elevation motor. The elevation motor in turn executes custom-made code saved into its controller.
- The sequence starts by checking whether or not the mirror's current position is one of the mechanical structure's elevation limits. During a continuous scan the elevation

Figure C.4: Diagram illustrating the devices and connections involved in the radar's real-time positioning subsystem.

motor performs rapid up-down movements through most of the structure's length and, since stopping requires a certain deceleration distance, starting the scan at a position too close to a mechanical limit can deteriorate the motor through collisions. This verification is performed by the microcontroller by checking the output of two inductive sensors (one in each limit), which provide a pulse when nearby metallic objects, such as the mirror, are detected. Photographs of the inductive sensors and the elevation limits on the mechanical structure are shown in Fig. C.5, and the mirror positions are shown in Fig. C.6. These sensors are critical to the positioning system, because they provide an external absolute reference for the motors' encoders, which are relative⁴.

Figure C.5: Photographs of (a) the inductive sensors and (b) the elevation limits in the mirror's mechanical structure.

- If the motors are in any position other than the upper limit (corresponding to an elevation angle of $\gamma = -6.8^\circ$), the microcontroller orders the elevation motor to move upwards towards this position. Once there, the motor is in a known starting point. Nonetheless, as mentioned before, mechanical limits are not recommended as a starting

⁴Meaning they only provide positioning information relative to the initial position at each system start.

Figure C.6: Chronological sequence of mirror positions during the vertical alignment. In the first step, the mirror begins the initialization in a random position. The motor is then ordered to move upwards until it reaches the upper structural limit, in the second step. In the final step, the motor has been ordered to move downwards a fixed number of steps to start the scanning in a mechanically safe position.

position. With this in mind, the motor is ordered to go down until its encoder outputs its first index signal. This signal corresponds to a given number of motor steps, and therefore allows knowing the distance travelled by the motor and establishing a new starting point.

- At this point in the sequence, the scan has a fixed vertical starting point (focusing on the uppermost ellipse of the scan pattern), but the initial position of the rotation motor has not yet been established. The rotation motor is now ordered to move at scan speed. At any given time during this rotation, the rotation motor passes through a third inductive sensor (another known location reference), and the microcontroller then starts keeping real-time count of the motor's steps through its encoder. The motor continues its rotation throughout the entire scan.
- When the rotation motor's pass through the inductive sensor is detected by the microcontroller, an Ethernet message is sent to the PC signalling that the mirror is now in a known position and the scan may start. The PC simultaneously starts signal generation through an USB connection to the DDS, sampling in the ADC, and the continuous scan. All three operations then start from a single PC-controlled time origin.
- From then on, and until the scan is stopped by the user, the elevation motor performs a continuous up-and-down movement, while the radar transmits its chirps, and the PC samples and processes the IF signal. The mirror's position is given by encoder outputs, processed in the PCB electronics, and sent to the PC via Ethernet.
- During the scan the ADC also samples an output from the elevation motor, which produces a pulse train indicating when the motor is moving or stopped, and serves as further synchronization. The PC, through C++ code, uses this signal to detect when the motor changes directions and how long each movement lasts. Examples of this signal for different scan speeds are presented in Fig. C.7.

Each encoder is able to distinguish between ~ 4000 different motor positions, allowing highly accurate localization. Since the motors are in continuous motion, after each frequency ramp the mirror will be located in a different position, and thus, encoder information is needed. Potentially, the microcontroller is able to sample, process, and transfer the data from the encoders in real-time as dictated by interrupt request (IRQ) pins (which are connected to the DDS's trigger signal, synchronizing ramp generation with mirror positioning).

Figure C.7: Examples of the elevation motor's output signal captured by the ADC, when the motor moves (a) at ~ 0.6 fps, and (b) at ~ 1.2 fps. When the pulse train is in its *high* state, the motor is in movement. Otherwise, the motor has stopped. Circles in red are the spots where the C++ software has detected movement starts or stops.

Mirror positions are extracted from the encoders through the use of 12-bit counters in the microcontroller's PCB. The pulse trains in the Quad Count channels (A/B) output from each motor's encoder are input into a counter, whose output is captured by GPIO pins in the microcontroller. When the positioning interrupt reaches the microcontroller, the voltage value on those pins is stored and sent to the PC. Meanwhile, the index signals (1 pulse per 360° for the rotation motor, and 1 pulse per 200 steps for the elevation motor), also captured by the microcontroller, are used as overflow control for the counters. A simple diagram of these connections is seen in Fig. C.8. Since the A/B channels are offset by 90° with respect to motor movement, the XOR gates allow doubling the positioning resolution by doubling the number of pulses per distance travelled.

Figure C.8: Block diagram showing the encoder output paths in the microcontroller PCB, which receives and transfers their information to the PC.

Knowing the number of pulses each encoder outputs per step (400 pulses per 200 steps, for the elevation motor) or rotation (2000 pulses per 360° , for the rotation motor), a relative position can be calculated from a starting point and the total number of steps travelled or turns performed. A pseudocode representation of the implemented algorithm is as follows:

Algorithm 1 Absolute mirror position extraction from encoder outputs.

```

1:  $\phiIncrement = 360/8000$   $\triangleright$  In  $^\circ$ /pulse. Divided by XOR gate and gear reductor
2:  $\phiAtIndex = [0, 90, 180, 270]$   $\triangleright$  1 index per quarter mirror turn, due to gear reductor
3:  $stepsAtIndex = [0, 200, 400, \dots, 1200]$   $\triangleright$  Measured values
4:  $distancePerStep = 0.0317$   $\triangleright$  In mm/step, from elevation motor datasheet
5:  $degreesPerDistance = 0.15/1.13$   $\triangleright$  In  $^\circ$ /mm
6: if mirror moving down then
7:    $sign = 1$ 
8: else
9:    $sign = -1$ 
10: end if
11:  $elevationSteps = stepsAtIndex[elevationIndex] + elevationPulsesCounted \cdot sign$ 
12:  $\gamma = \gamma_{start} + elevationSteps \cdot distancePerStep \cdot degreesPerDistance$ 
13:  $\phi = \phi_{start} + \phiAtIndex[rotationIndex] + \phiIncrement \cdot rotationPulsesCounted$ 

```

In practice, some restrictions to this scheme have emerged, due to microcontroller limitations and the low throughput of the Ethernet connection. The result of timing 100 positioning data packets over the two different transmission protocols available is shown in Fig. C.9. As expected User Datagram Protocol (UDP) achieves lower times than Transmission Control Protocol (TCP) due to the removal of error-correcting checks and guarantees, but in both cases transmission lasts longer than the duration of a chirp signal (150 μ s). During a continuous scan this ends up either blocking the IRQ service through multiple interruptions or interrupting an ongoing message, where both cases eventually cause the loss of packets.

Figure C.9: Histograms corresponding to 100 positioning data packets sent over (a) TCP and (b) UDP through the Ethernet connection between the PCB and the PC.

The result of implementing the real-time positioning mechanism is shown in Fig. C.10, where the radar has scanned the FoV with two configurations: the offline-interpolated positioning used before this thesis (in red), and real-time 1-packet-per-ramp positioning (in blue). As can be seen, the FoV is slightly extended upwards and downwards in the new configuration, spots positions are non-uniform, and entire ellipses are missing from the scan. The effects can be seen more clearly in the zoomed section. Individual spots are missing in all ellipses, meaning single data packets have been lost or unsent (as explained by the transmission histograms). The total number of spots for the offline data is 7800, while that figure goes down 37.91% (to 4843 spots) in the encoder-based information. Also, entire sections of

the scan (composed of hundreds of spots) are empty.

Figure C.10: Comparison of the interpolated offline-positioning (in red) and the real-time 1-packet-per-ramp positioning (in blue) techniques. In the zoomed view, some of the limitations of the real-time positioning are better seen. Individual spots are lost due to microcontroller clogging, and entire sections are empty due to motor speed misestimation.

In the final implementation, to avoid missing spots in any area of the FoV mirror position data are sent only when the elevation motor changes directions and thus starts a new image, instead of after every chirp generation. This way, the first coordinate in each image is initially synchronized with the mirror's position. The rest of the coordinates are then extrapolated from that initial position assuming constant motor speed, and using the polynomial functions presented in Section 2.2.1. A diagram of this scheme is presented in Fig. C.11.

Figure C.11: Chronological sequence of events in a typical continuous scan. When the mirror reaches the upper or lower scanning limits, a data packet is sent to the PC containing the initial position information for the next scan. While the elevation motor is moving up or down, no data packets are sent (positioning is done via interpolation from the initial position)

From a throughput point of view, sending just one packet per scan generates the least strain possible on the positioning system. Alternative methods with higher data rates can also be implemented, such as sending one packet per every n chirp signals. This can be achieved by setting a digital counter on the microcontroller IRQ pin, allowing interrupts only every n calls. Under this mechanism, the radar would have real-time positioning information every n spots, in between which (x,y) interpolation would take place. The larger the n value, the longer the PC must interpolate mirror positions, and the less accurate those positions will be. Additionally, setting an n value too large can cause skips in elevation angle estimation. On the other hand, a very low n value will create overlapping of microcontroller interruptions and produce the same packet loss also seen in Fig. C.10.

Examples of the 1-packet-per- n -ramps scheme, for n values of 10, 100 and 1000, can be seen in Fig. C.12.

Figure C.12: Comparison of the interpolated offline-positioning (in red) and the real-time 1-packet-per- n -ramps positioning (in blue) techniques. The n value is 10 in (a), 100 in (b), and 1000 in (c).

With an n value of just 10, a noticeable improvement in scan continuity can be noticed when compared with Fig. C.10. Nonetheless, several spots are still missing throughout the scan. Incrementing the n value to 100, no discernible improvement is seen in continuity, and multiple spots are still lost. Finally, when the n value is risen to 1000, the PC only needs to receive approximately 8 positioning packets, which reduces the number of missing spots to zero (since the vast majority of positions are now interpolated). Although an n value of 1000 seems like an appropriate factor, the 1-packet-per-scan scheme has been the implemented choice as a safeguard to ensure the lowest packet losses possible.

The 1-packet-per-scan mechanism can be misinterpreted as being the same as a scheme where only one packet is sent during the initialization process and then a continuous interpolation is performed for the duration of the scan. This can be erroneously thought of as viable alternative due to the fixed nature of the number of steps that the elevation motor moves in every up-and-down sequence, and its stable speed. To showcase this nature, the data in Fig. C.13 have been collected, where the movement of the elevation motor is timed during a continuous scan consisting of 100 images, for varying speeds. Unfortunately, sending only 1 packet during set-up cannot be done without adding a rotational error (angle ϕ) to all images after the first one, since the rotation motor does not stop its movement during the scan and its high speed (8000 rpm in a typical scan) requires a very fine timing estimation. At least one measurement of the initial ϕ angle of the mirror is required for each image.

The stable nature of the rotation motor has also been studied. Since the motor is ordered to rotate at a fixed speed throughout the complete continuous scan and, as seen in Fig. C.13, up-and-down movements have little time variation, the initial ϕ angle registered for each individual image takes on discrete values. For an initial ϕ_0 during the radar's initialization process, the starting rotation angle in each of the following images (ϕ_i) will be

$$\phi_i = [\phi_0 + (6 \cdot \theta) \cdot (\sum_{n=1}^{i-1} (t_{movement_n} + t_{stop_n}))] \pmod{360} \quad (C.1)$$

where θ is the rotation speed (in *rpm*) and $t_{movement_n}$ and t_{stop_n} the elevation movement and stop times, respectively, for image number n . This disposition has been tested for multiple rotation speeds during a continuous scan, and the results are presented in Fig. C.14. As

Figure C.13: Registered times for two different elevation events (movement and stops) during a 100-image continuous scan, for varying elevation motor speeds. The data around 0.1 s (equal for all speeds) correspond to the duration of stops in the upper or lower limit when an individual scan is finished, and it is limited by the motor controller and its code. Little variation in time required can be appreciated during movement at any speed.

expected, as the speed increases the variation in recorded angles also grows larger, since the slight changes in elevation and stop times create bigger rotation offsets.

Figure C.14: Histograms of registered initial rotation angles for 100 consecutive images during a continuous scan, for varying motor speed.

The histograms of recorded angles in a continuous scan can also be represented in a chronological format, where the angles sent in each UDP packet are ordered according to their arrival. The resulting graphs are presented in Fig. C.15 for the rotation motor and Fig. C.16 for the elevation motor.

In these representations the discrete nature of the initial angles is evident: because the up-and-down movement is performed with a fixed number of steps, and fixed rotation speed, the angles alternate between a finite set of values. Within the rotation sequences, as expected, the angles at the highest speed (8000 rpm) present the larger variation, due to the bigger relative error from the spread of the elevation's duration. Elevation sequences only alternate between two possible angles (top and bottom), and remain almost identical in different scans.

Figure C.15: Registered initial rotation angles in chronological order for 33 consecutive images during a continuous scan, for varying motor speed.

Figure C.16: Registered initial elevation angles in chronological order for 33 consecutive images during a continuous scan, for varying motor speed.

C.2 Bandpass Sampling

Another loss-of-data problem encountered was due to ADC sampling. The current ADC installed⁵ is not able to continuously sample the IF signal over an undefined time period without overflowing, and thereby losing samples. During a real-time scan this means that the PC receives incomplete images and the synchronization mechanism is broken (the following images, even if complete, will not be correctly positioned). A simple verification of this problem was performed by feeding the ADC with a sinusoidal signal generator output. The result, for two sampling scenarios, may be seen in Fig. C.17. In Figs. C.17(a) and (b) the largest sample losses correspond to 1.94 s and 2.56 s of capture time, respectively. This means that ~ 13000 and ~ 17000 chirps⁶, respectively, would be lost. In both cases, the loss appears after around 80 s of capture, which would limit continuous scanning to less than two minutes of video.

Figure C.17: Examples of ADC buffer overflow during the continuous sampling of a sinusoidal signal. The red dashed rectangles indicate where samples have been lost. In (a) the sampling frequency is 40 MS/s, and only 1 ADC channel is being captured. In (b) the sampling frequency is 25 MS/s, and 2 ADC channels are being captured.

For the real-time algorithm designed, two ADC channels (one for the IF signal and one for the elevation motor) are necessary. The large throughput that this requirement produces is the cause for buffer overflow. For the ADC used, the connection to the PC is via PCI 1.0. The ADC supports two variants of transfer speed for this interface: a 32-bit/33-MHz version (at 133 MB/s), and a 32-bit/66-MHz (at 266 MS/s). In bps, their corresponding throughputs (from ADC buffer to PC memory) would be:

$$T_{33 \text{ MHz}} = 133 \frac{\text{MB}}{\text{s}} \cdot 8 \frac{\text{bits}}{\text{Byte}} = 1064 \text{ Mbps} \quad (\text{C.2})$$

$$T_{66 \text{ MHz}} = 266 \frac{\text{MB}}{\text{s}} \cdot 8 \frac{\text{bits}}{\text{Byte}} = 2128 \text{ Mbps} \quad (\text{C.3})$$

Additionally, for the usual sampling frequency (40 MS/s) and two ADC channels, the data capture throughput (from ADC input to ADC buffer) would be:

$$T_{40 \text{ MS/s}, 2 \text{ channel}} = 40 \frac{\text{MS}}{\text{s}} \cdot 2 \text{ channels} \cdot 16 \frac{\text{bits}}{\text{Sample}} = 1280 \text{ Mbps} \quad (\text{C.4})$$

⁵Model PCI-9846 from ADLINK Technology.

⁶Both amounts corresponding to more than two full images each.

Unfortunately, for the motherboard used, the 32-bit/66-MHz is not an option due to the Platform Controller Hub (PCH) that governs the PCI transfers in the PC, which only permits 33 MHz transfers through PCI 1.0. As can be seen, $T_{40MS/s, 2channel} > T_{33MHz}$ which means that overflowing will eventually occur. As illustration, the theoretical required time for various data transfer interfaces to send a given number of bytes can be seen in Fig. C.18. The resulting curves are then compared against the time needed for the ADC used to generate that total data size when capturing two channels at different sampling rates. The results confirm that for 40 MS/s sampling, 32-bit 33 MHz PCI is not a viable option. At their theoretical maximum rates, all other interfaces can support this speed.

Figure C.18: Data transfer for various interfaces compared against varying sampling rates, as total data size increases and 2 ADC channels are captured.

To overcome this obstacle, several possible solutions are discussed below:

- **Installation of a new motherboard.** In order to maintain the programming structure of the developed algorithm, the motherboard would need to still use an Intel processor. Intel's i7 series, and above, do not allow 66 MHz transfers through PCI 1.0 connection. A reduction to a less powerful processor is also not viable due to possible limitations to the real-time signal and image processing code.
- **Installation of a new ADC.** The use of a new ADC with a newer PC interface could solve the issue. For instance, an ADC with PCIe 2.0 x1 support can use a throughput of 250 MB/s, or:

$$T_{PCIe\ 2.0\ x1} = 500 \frac{MB}{s} \cdot 8 \frac{bits}{Byte} = 4000\ Mbps \quad (C.5)$$

For instance, ADLINK's own PCIe-9852⁷ has PCIe 2.0 x4 support, making the available throughput four times $T_{PCIe\ 2.0\ x1}$. Even though a new ADC from the same manufacturer enables simple substitution, without implementing multiple changes in the algorithm, the price of this solution⁸ goes against the cost-efficient nature of the radar system.

- **PCI 1.0 to PCIe adaptor.** The current motherboard installed has an empty PCIe 2.0 x1 slot which could be used. Even though no ADCs with a single PCIe slot could be found, an adaptor can be used to allow the current ADC to connect with the PCIe

⁷Which includes 2 input channels, and up to 200 MS/s and 14 bits per sample.

⁸More than \$4000 for the ADC mentioned.

2.0 x1 slot. Since PCIe buses are directly controlled by the PC's processor, this would allow bypassing the PCH's 33 MHz limitation and use the current ADC's 32-bit/66 MHz interface. Although potentially possible, this option has not been implemented because all adaptors found use a 64 bit interface for their PCI connection (which is not compatible with the current ADC).

- **Lowering sampling frequency.** Since throughput is directly proportional to sampling frequency, lowering the number of samples captured per second would reduce the limitation. On the other hand, sampling frequency is also proportional to the maximum range detected without aliasing, thereby limiting the standoff distance. The radar's focus is located at 8 m, although due to propagation through the reflector system, the registered distance is usually around 10 m. Substituted into (1.18), this results in a 12 MHz beat frequency. Thus, at least 24 MS/s are needed to sample the IF signal without aliasing, according to the Nyquist criterion. Since not only detections from the target's surface are desired, but also its various range layers, an additional range margin must be established. For instance, for the 30 cm used during range windowing in the signal processing algorithm, the sampling frequency would need to be at least 24.72 MS/s. Unfortunately, as illustrated in Fig. C.17(b), even when sampling as low as 25 MS/s the ADC still overflows. With a 25 MS/s sampling frequency ADC overflow should not occur, as indicated by the associated throughput

$$T_{25MS/s, 2\ channel} = 25 \frac{MS}{s} \cdot 2\ channels \cdot 16 \frac{bits}{Sample} = 800\ Mbps \quad (C.6)$$

which is lower than $T_{33\ MHz}$. This is due to additional overhead in PCI connections when performing large data transfers, which reduces the actual throughput below the theoretical maximum. Through the same measurements as in Fig. C.17, it has been found that the ADC can support 5 MS/s sampling without losses, as shown for a four minute capture in Fig. C.19.

Figure C.19: Example of uninterrupted capture with the ADC during continuous sampling of a sinusoidal signal. The sampling frequency is 5 MS/s and 2 ADC channels are being captured.

In a conventional lowpass sampling scheme, a 5 MS/s sampling frequency would only allow distances up to 2.08 m, which is far below the radar's focusing range. To solve this difficulty, bandpass sampling has been implemented. The bandpass sampling scheme is summarized in Fig. C.20. To avoid aliasing, an analog bandpass filter has been added to the IF signal path before the ADC (Annex B.2), eliminating all detections closer than 7.5 m and farther than 12.5 m.

Figure C.20: Diagram illustrating the bandpass sampling and filtering process over the received signal's spectrum, used to enable continuous scanning with the current ADC.

Annex D

Modifications to a 122 GHz Commercial Radar

The active sensor described in Section 3.1 is based on a 122 GHz FMCW radar kit fabricated by Silicon Radar GmbH [453], shown in Fig. D.1, and priced below 1000 €. The kit is composed of three stacked boards, with the top two, as shown in the photograph, designed by Silicon Radar. The lower board, in white and blue, is a COTS microcontroller tasked with system configuration, signal processing, and data transmission to a PC. In the middle, a green offset PCB includes baseband electronics (amplifiers, filters, etc...). Lastly, the top board, directly below the lens, houses a packaged MMIC with the radar front end. The kit also includes a plastic lens, which provides overall directivity (thus allowing farther range detection), and PC software for browser-based display of the received spectrum. FMCW operation is controlled through the provided software, with a bandwidth of more than 15 GHz and sweep times from 12 μ s to 18 ms. Research surrounding this sensor may be found throughout the literature [204].

Figure D.1: Photograph of Silicon Radar's 122 GHz FMCW radar kit.

The radar's summarized architecture may be seen in Fig. D.2 (specific values have been omitted to protect intellectual property), and is further detailed in the referenced publications. Signal generation is achieved via a PLL synchronized with a VCO at 60 GHz. After doubling the signal's frequency, transmission is performed by a dipole antenna. Reflected radiation is received by another such antenna, both placed directly on the chip [94]. Two mixers in the MMIC provide IQ signals, amplified and filtered in the baseband PCB. These signals are sampled and processed by the microcontroller, which later sends the resulting data to a PC.

Figure D.2: Initial (as given by Silicon Radar’s kit) and modified radar architectures. Items in red have been removed, items in green have been added, and items in black remain untouched.

Initially, the radar’s transmission was limited to intervals managed and limited by the microcontroller. The voltage ramps attacking the VCO were generated in clusters, after which generation would stop to allow the microcontroller time used for on-board signal processing and USB-to-PC data transmission. The clusters of FMCW sweeps, with off-time between ramps, were separated by additional off-time (due to repetitive configuration of the transmitted waveform), as shown in Fig. D.3.

Continuous chirp transmission is necessary for the radar applications presented in this work, such as standoff imaging or vital sign monitoring. In order to enable continuous FMCW operation, a number of modifications were performed on the original sensor:

- Removal of Previous Microcontroller:** Since most restrictions stemmed from the use of the microcontroller board, a new PCB was designed to substitute it, which may be seen in Fig. D.4(a). This new control board includes an Arduino Nano for FMCW configuration, but signal sampling and processing are now done off-board. The new microcontroller only configures the desired waveform once (at system start), unlike the previous device which sent its configuration after every individual chirp (thereby causing intra-cluster off-time). Additionally, the power source of the microcontroller (USB port) can be used independent from the PLL’s (which can be through a voltage-regulated jack connector or dedicated pins), providing a cleaner source for waveform generation.
- Access to Signal Synthesizer:** By removing the microcontroller, direct access to the PLL’s synthesizer and phase-frequency detector [Analog Devices ADF4169; shown in Fig. D.4(b)] is now available without the use of the original software. This approach, via an Arduino board with custom-made code, enables a wider range of waveform configurations than previously available, including finer bandwidth selection, a larger selection of chirp times, and the introduction of properly-aligned flyback times for frequency sweeps. Furthermore, the synthesizer’s end-of-ramp signal is now used as a

trigger signal for the ADC, achieving sweep-synchronized sampling. The modifications to access this trigger signal are shown in Fig. D.4(a) (the addition of a comparator) and Fig. D.4(c) (the removal of a transistor).

- Sampling and Processing:** These functions are now performed on a PC. Sampling is done on a PCI-connected ADC¹, enabling data transmission at much faster rates than the previous USB-based communication. Additionally, the new ADC also has better resolution and a higher sampling frequency, producing better SNR and enabling higher beat frequencies. By redirecting a sample of the PLL's reference oscillator to the clock input of the ADC, signal generation and sampling are now locked (a frequency divider in this path allows choosing different sampling frequencies from the same reference) thus reducing phase uncertainty when sampling. A 1:4 buffer is also included before the frequency divider, which would allow locking three other sensors to the same PLL reference (allowing, for instance, synchronized node networks or array-based imaging). Furthermore, processing is now performed on the PC, with custom-made code in C++ and/or MATLAB.

Figure D.3: Voltage ramps input into the MMIC's VCO for FMCW operation in the modified (in blue) and original sensor (in red), showing both types of off-time preventing continuous transmission in the unmodified case. In the larger graph the blue data shows that transmission is now continuous, while the zoomed-in view confirms that it is composed of voltage ramps of the same amplitude as before the modifications and with faster sweeps.

The PLL's reference used as clock signal for the new ADC is accessed by tapping into the path between the 75 MHz crystal oscillator² and the synthesizer, as shown in Fig. D.4b. This signal is then fed to the new control board via SMA connector and cable. In the control board the signal passes through a 1:4 buffer, where one of the outputs is directed to a frequency divider. The buffer and the divider are required in order to condition the signal to the sampling characteristics of the new ADC: the frequency must be lower than 40 MHz, the signal AC-coupled, and the maximum amplitude between 1 V_{pp} and 2 V_{pp}. A comparison displaying the clock signal at different stages can be seen in Fig. D.5, where the divider has been set to a 1:4 ratio (outputting 18.75 MHz). Notice how the clock signal after the divider is indeed four times slower than before, and its amplitude has been reduced to ADC-safe levels. It can also be noticed that the initial clock signal is a sine wave (typical in RF components), while the divided signal is a pulse train (more common in digital circuitry).

¹Adlink PCI-9846

²Abracon ASDMB-75-LY-T

Figure D.4: (a) Photograph of the top side of the new control board. (b) Photograph of the bottom side of Silicon Radar's modified RF board. (c) Photograph of the top side of Silicon Radar's modified RF board.

Figure D.5: (a) Comparison of the clock signal at different places in the boards. (b) Zoom of the final clock signal used for sampling in the new ADC.

The frequency response of the 75 MHz reference oscillator, along with its phase noise performance, measured at the soldered SMA output may be seen in Fig. D.6. Similarly, Fig. D.7 shows the same measurements for the outputs of the frequency divider introduced in the modified PCB. As can be seen from the spectrum, the frequency divider introduces powerful x2 harmonics for any selected division ratio. On the other hand, these harmonics are as narrow as the desired frequency tone, so no interference is detected. This is further seen in the phase noise data, by the comparison with its input (from the reference oscillator), where the divider does not alter the phase noise performance for any of the division ratios available. In addition, the measured phase noise is above the spectrum analyzer's displayed average noise level (DANL), confirming that its source is the radar and not the measurement equipment.

Figure D.6: (a) Spectrum and (b) phase noise performance of the 75 MHz reference oscillator.

Figure D.7: (a) Spectrum and (b) phase noise performance of the frequency divider.

The comparator shown in Fig. D.4(a) is necessary to condition the trigger signal to the new ADC. The synthesizer's end-of-ramp signal contains two pulses per ramp: one for the end of the rising frequency sweep, and one for the decreasing sweep (flyback). These pulses, which can be seen in Fig. D.8(a), are similar in duration ($\sim 0.5 \mu\text{s}$) but very different in amplitude (0.3 V for the first, and 1.6 V for the second). Only the second pulse is useful

in this case, since the flyback samples must also be taken into account when processing. By sampling the received signal in synchrony with this pulse the entire ramp (upwards and downwards sweeps) is sampled. To suppress the first pulse, before reaching the ADC, a comparator is included in the control board. This component compares the end-of-ramp signal with a fixed threshold (adjustable with a potentiometer) and outputs only the second pulse, without major modifications to its shape [as seen in Fig. D.8(b)].

Figure D.8: (a) Comparison of the end-of-ramp signal from the RF board's synthesizer, before and after passing through a comparator. (b) Zoom of the pulse associated to the downwards frequency sweep, before and after passing through a comparator.

The new ADC performance may be seen in a comparison with the previous configuration in Fig. D.9. Due to the use of a new ADC, the maximum sampling frequency, and thus maximum detectable range, has been increased from 5 MS/s to up to 40 MS/s. The dc component has been reduced due to the addition of a RC high-pass filter in the IQ paths of the new control board. The noise floor in the working range has also experienced a reduction, of approximately 20 dB, which can be explained by the elimination of short-range leakage (via locking the ADC's clock with a sample of the synthesizer's oscillator, and filtering the dc component) [206].

As opposed to the initial sensor, in which transmission was interrupted, with continuous generation in the modified radar more average power is transmitted. This results in higher heat dissipation in the MMIC. To analyze this effect, a measurement of the evolution of the MMIC's temperature was made, using an IR thermometer, from cold start to stabilization during continuous transmission. As the results in Fig. D.10 show, the temperature quickly rises to 55 °C, after which it remains stable. Since the antennas' substrate seems to be Rogers Ultralan 3850 [204], which presents a higher thermal resistance than this value, no temperature related problems are expected.

To solve some of the limitations of the modified radar, such as its size (which creates clutter and blanking during imaging scenarios) and transmitted phase noise (which decreases dynamic range), an updated system has been designed. A visual comparison between the two radars is presented in Fig. D.11. The new board only has a 50x50 mm² area, greatly reducing the size from the previous version. This is mainly thanks to an increase in component density and the elimination of Silicon Radar's baseband board, whose signal conditioning functions are now performed by a single board. Additionally, the Arduino microcontroller has been substituted for a less-powerful, but also less bulky, alternative. Although this update has been performed after most of this thesis' work was completed, very promising improvements

Figure D.9: Comparison of the noise floor and sampling frequency for the previous and new configuration of the active imaging sensor. A target is placed at 2 m, sweep and FFT time are $600 \mu\text{s}$, and 6.3 GHz of bandwidth are transmitted

Figure D.10: Evolution of MMIC temperature during continuous transmission. Temperature becomes stable after approximately 10 minutes of uninterrupted operation.

in the radar's performance have been accomplished, which would warrant a renovation of some of the measurements presented in this document.

Figure D.11: Photograph comparing the modified 122 GHz radar board developed during this thesis (left) and its updated version (right).

Annex E

Equipment and Material for Data Acquisition

The documents included in this appendix are the specifications of the various equipment and materials used during the data acquisition measurements of Section 5.1. The documents are organized as follows:

- Datasheet from Menlo Systems for the TPX50 and TPX200 THz focusing lenses used in Sections 5.1 and Annex F.2. Both lenses are made of TPX (which is the Mitsui Chemicals trademarked name for polymethylpentene); a material that presents an absorption coefficient of $\alpha = 0.08 \frac{NP}{cm}$ and a refractive index of $n = 1.46$ [346], causing a reflectivity of $R_o = 0.035$ from

$$R_o = \left| \frac{n_{air} - n_{TPX}}{n_{air} + n_{TPX}} \right|^2 \quad (E.1)$$

where $n_{air} = 1$. Their thickness is ~ 3 mm, which considering α would produce 0.21 dB transmissivity losses through

$$L_{transmissivity}(dB) = \alpha \cdot \frac{20}{\log_n(10)} \cdot 0.3 \text{ cm} \quad (E.2)$$

and a 14 dB-attenuated ($A_{reflection}$) reflected signal compared to the transmitted one due to

$$A_{reflection}(dB) = |10 \cdot \log_{10}(R_o)| \quad (E.3)$$

- Reference sheet from Phemsa S.L. for the HDPE material machined into the data acquisition guides used in Section 5.1.
- CAD rendering and photographs of the guide designs used in Section 5.1, with dimensions in mm.

TPX50/100/200

Lenses for THz Radiation

TPX50

TPX100/200

Our TPX lens series is optimized for collimation and focusing in THz optics. The TPX50 is a plano-convex lens with a effective focal length (EFL) of 50 mm, designed for optimal collimation of highly divergent THz radiation. TPX100 and TPX200 are aspheric lenses with 100 mm and 200 mm effective focal length, respectively, optimized for focusing the THz radiation for increased working distance in the terahertz path.

The TPX lenses are mounted in a 1.5" lens holder with a clear aperture of 35.5 mm.

TECHNICAL DRAWING TPX50

MenloSystems

KEY SPECIFICATIONS

- Broadband THz Transmission
- Light-Weight TPX Material, Transparent for 780 nm and 1560 nm
- TPX50: EFL 50 mm
- TPX100: EFL 100 mm
- TPX200: EFL 200 mm

APPLICATIONS

- THz Collimation and Focusing
- Broadband THz Spectroscopy

FEATURES

- Broadband Transparency in THz Regime
- Light-Weight TPX Material for Long Term Stability
- Mounted with Free Aperture of 35.5 mm

TPX50/100/200

MenloSystems

Lenses for THz Radiation

SPECIFICATIONS	TPX50	TPX100	TPX200
Wavelength	Broadband THz range	Broadband THz range	Broadband THz range
Material	TPX polymer	TPX polymer	TPX polymer
Lens Geometry	plano convex, aspheric	plano convex, aspheric	plano convex, aspheric
EFL Effective Focal Length	50 mm	100 mm	200 mm
Lens Dimensions	1.5"/38 mm diameter	38 mm diameter	38 mm diameter

MOUNTING EXAMPLE*

*delivered without rail and post holders

ORDERING INFORMATION

Product Code	TPX50	TPX100	TPX200
Price	EUR 220	EUR 220	EUR 220

For OEM quantities please call for pricing. Prices and Specifications are subject to change without notice.

MenloSystems

Menlo Systems GmbH
T+49 89 189 166 0
sales@menlosystems.com

Menlo Systems, Inc.
T+1 973 300 4490
ussales@menlosystems.com

Thorlabs, Inc.
T+1 973 579 7227
sales@thorlabs.com

www.menlosystems.com

www.frequencycomb.com

D-TPX50TPX100-EN 05/10/16

PHEMSA

Plásticos, herramientas y metales.

HDPE Reference Sheet

Polietileno

Características generales

Polietileno de alto peso molecular. Muy rígido, con un bajo coeficiente de fricción y una buena resistencia a la abrasión. Es apropiado para usar en contacto con productos comestibles y usado en plantas químicas. Puede ser unido con técnicas de soldadura mediante gas caliente y puede ser manufacturado con facilidad.

Propiedades

- Peso específico: 0,96 kg./dm³
- Absorción de agua a saturación: 0,02 %
- Punto de fusión: 135 °C
- Temperatura de trabajo "mínima": -50 °C
- Temperatura de trabajo "máxima": 80 °C
- Resistencia a la tracción: ASTM D-638 26 MPa
- Resistencia a la compresión: ASTM D-695 28 MPa
- Resistencia a la flexión: ASTM D-790 30 MPa
- Módulo de elasticidad: 10.000 kp/cm²
- Alargamiento a la rotura: 600 %
- Constante dieléctrica: 105Hz2,3
- Rigidez dieléctrica: 150 Kv/mm
- Coeficiente de fricción P=0,05 N/mm²; V=0,6 m/s: 0,29
- Resistencia específica: 1017Wcm
- Color: Blanco y Negro

First Guide Design

Second Guide Design

Third Guide Design

Annex F

Material Analysis with FMCW Radars

In this annex, a simple analysis of the absorption and reflection characteristics of some common plastics, with a special emphasis given to plastics used as base material for 3-D printing, is presented. The intended application for the analysis would be providing material characterization useful to distinguish between different kinds of plastics by FMCW scanning, which could directly be exploited, for instance, in recycling plants to efficiently sort waste [454]. For this thesis, the analysis has been limited to measuring the losses from different plastics due to material density, inner structure, orientation, and dye. In the future, permittivity and permeability extraction techniques used for material identification could be tested [455], benefitting from the attenuation analysis here presented.

In addition, this appendix also deals with material characterization measurements applied to common clothing fabrics, presented in Section F.2 [23, 456]. The motive behind this study is to begin the development of a clothing database for the security screening systems, which could be used to compensate the signal attenuation and image distortion created by different clothes when scanning a person [457, 458]. As was the case for Section F.1, textile analysis in this study has been limited to measurements of the reflection losses from various fabrics and orientations.

Both analyses use measurement set-ups and approaches established in Section 5.1, for both the 300 GHz and the 122 GHz radars. For the 122 GHz device, as was the case for the measurements in Section 5.1, only results from the MMIC configuration with the largest bandwidth (**Configuration 3**, with 15.1 GHz of transmitted bandwidth) are presented¹.

F.1 Analysis of Absorption in Plastics

The measurement set-up used in Section 5.1 has been applied to a study of mm-wave absorption by common plastic materials at 300 GHz and 122 GHz. In this analysis several variables have been taken into account, such as material type, density, and inner structure. The materials tested are HDPE and Polyactic Acid (PLA), another polymer commonly used as 3-D printing filament. For both materials, two different densities have been used (machined from larger blocks in the case of the HDPE, and 3-D printed for the PLA). In addition, the PLA tests include variation regarding the dye color used for the filament. This last variable is motivated by an obstacle of the waste management industry: traditional plastic-separation

¹The two remaining configurations have also been evaluated, but have shown worse performance.

systems (based on optical sensors) in recycling facilities are not capable of accurately distinguishing plastic types when the materials are painted black. This situation has induced the rise of several research projects dedicated to solving the separation problem via mm-wave sensors (which unlike cameras are impervious to lighting and color changes) [459, 460], along with other devices [461], including a fully-operational sorting system [462] (shown in Fig. F.1).

Figure F.1: Photographs of the blackValue[®] sorting system, developed by the Fraunhofer Institute for High Frequency Physics and Radar Techniques. Image reproduced from [462].

The solid HDPE slabs used are both 50x200x300 mm³ in size, while their densities are $0.93 \frac{g}{cm^3}$ and $0.96 \frac{g}{cm^3}$ for the first and second pieces, respectively. The PLA designs are 60x40x100 mm³ in size, printed with the same resin², but at 10% and 50% density for the first and second pieces, respectively. The 3-D printed blocks are built with a 45° rectilinear infill pattern shown in Fig. F.2, with solid faces in four of their sides (leaving two of them open to see the infill pattern). For all measurements of the pieces in this section, the signal processing steps described in Section 5.1 are followed.

Figure F.2: Photographs of the PLA pieces considered for evaluation. (a) Black-colored PLA of 50% (top, *PLA1*) and 10% (bottom, *PLA2*) rectilinear fill pattern, (b) Red-colored PLA of 50% (top, *PLA1*) and 10% (bottom, *PLA2*) rectilinear fill pattern.

A first comparison among the two HDPE pieces, with the spectrum from a reference

²Detailed in an official datasheet from the manufacturer (Mundo Reader S.L.) at the end of this section.

metallic sheet included, may be seen in Fig. F.3 at 300 GHz and in Fig. F.3(b) at 122 GHz. At 300 GHz, both pieces of plastic return the expected levels of power from Section 5.1. It is interesting to note that the difference in density does not seem to produce a substantial change in the reflected power from any of the pieces' surfaces, since both peaks present similar levels among the two pieces. For the front surface the returned power is -43.01 dBm for HDPE1 and -42.51 dBm for HDPE2, while for the back surface these figures are -44.34 dBm and -44.23 dBm, respectively. The difference in these values is so small that they can be completely attributed to set-up uncertainty. This would indicate that the 300 GHz radar is not able to distinguish HDPE blocks of those two densities, based solely on the amplitude of their reflections. For pieces with density variation, one would expect a difference in the returned power from the back peak (as absorption through the material is increased) for a constant front-peak reflection. Reflections within the pieces (between 33 cm and 40 cm), although different, also register similar power. For the 122 GHz sensor, both HDPE pieces register practically equal returned power, indicating that the radar is also not able to distinguish between the two material densities based solely on amplitude information.

Figure F.3: Spectrum comparison at (a) 300 GHz and (b) 122 GHz (**Configuration 3**) of a reference metallic sheet and the two HDPE pieces considered.

The equivalent comparison for the two black PLA pieces is presented in Fig. F.4 at 300 GHz and 122 GHz. In the 300 GHz case, several differences can be noted with regard to the HDPE measurement. First of all, the power variations from the front-surface peak are now much more considerable. For the PLA1 piece, the reflected power from this surface is -43.49 dBm and -53.55 dBm, when illuminating the solid and open face, respectively. Meanwhile, for the PLA2 design, these values are -39.26 dBm and -58.22 dBm, respectively. The large differences between radiating an open or solid side (~ 10 dB for PLA1 and ~ 20 dB for PLA2) indicate that the radar would be capable of detecting when it is facing each side of the plastic. The reflection difference between faces is simply due to the density of material in each side: an open side presents empty holes from the rectilinear infill, which produce no reflection. In this regard, PLA2 presents a larger difference between side because of its lower infill density, which creates bigger holes in the pattern. On the other hand, the ~ 4 dB difference in returned power from the solid face between the two densities is likely due to set-up error. Also of interest is the wide peak from ~ 34 cm to ~ 37 cm for the PLA2 piece when illuminated on its solid face. These are joined reflections from the air-plastic interfaces of the infill pattern (which have smaller separation than the range resolution and are thus not properly differentiated), which are not present when illuminating the open face of the PLA1 design due to its denser infill (which effectively acts as solid block of PLA, without air interfaces). Through the existence of these internal peaks the radar could distinguish

between PLA designs of different densities, provided that the difference is as high as tested here. In addition to these considerations, it can be noted that no consistent back-surface peak can be detected for any scenario, even though the pieces' sizes are larger than the range resolution, evidencing that with PLA objects the radar is subject to stronger absorptions than with HDPE.

Figure F.4: Spectrum comparison at (a) 300 GHz and (b) 122 GHz (**Configuration 3**) of a reference metallic sheet and the two black PLA pieces considered, when illuminating them on a solid side or an open side (exposing the piece's infill pattern).

When the comparison is done for the red PLA pieces, as in Fig. F.5, additional considerations can be extracted. At 300 GHz, first, the power difference in the front peak between illuminating the open or solid face is still present, being the returned power -41.40 dBm for the solid face (for both densities) and -49.37 dBm and -53 dBm for the open face in the high- and low-density designs, respectively. This means that the power difference is still as large as ~ 8 dB, allowing for a separation threshold. For these two pieces also, no back peak is detected. Additionally, no substantial differences compared to the black PLA designs can be noted, indicating that the radar is indifferent to the coloring of the filament used. Side-by-side comparisons between the two PLA colors can be seen in Fig. F.6.

In the 122 GHz case, as for the 300 GHz radar, a large returned power difference can be noted when illuminating the open side versus the solid face of the PLA blocks. This is further evidenced when separating the spectrums according to the side that is facing the sensor, such as in Fig. F.7. Although slight differences in returned power for the different PLA densities are consistent in the graphs (the returned power for the first peak seems to always be higher for the PLA1 design), it cannot be fully determined whether the radar is able to distinguish between the two densities simply from reflection levels. On the other hand, focusing on Fig. F.7(a), it can be noticed that for the PLA2 design (when illuminating the solid side) the back-surface peak can be detected. This was not possible for the 300 GHz sensor, and enables the use of 3-D printing materials for the guides in the data-acquisition technique of Section 5.1. In addition, no substantial variations between the spectrums of the black and red PLA at 122 GHz can be detected, which would further prove that the mm-wave sensors are indifferent to plastic coloring (contrary to optical sensors).

When compared with the HDPE slabs, such as in Fig. F.8 at 300 GHz, the PLA pieces show similar returned power for their front peak when illuminating the solid face, while a larger variation is present when the radar radiates the open side of the PLA designs. In addition, as previously mentioned, the presence of a back-surface peak for the HDPE blocks is not duplicated for the PLA pieces.

Figure F.5: Spectrum comparison at (a) 300 GHz and (b) 122 GHz (**Configuration 3**) of a reference metallic sheet and the two red PLA pieces considered, when illuminating them on a solid side or an open side (exposing the piece's infill pattern).

In conclusion, the measurements performed indicate that in a plastic-sorting application the 300 GHz radar would be capable of detecting when different plastic materials (such as HDPE and PLA) are radiated through the existence of back-surface peaks. Furthermore, if the material's density affects its inner structure (creating air-plastic interfaces), differences in density can also be detected. On the other hand, the dye used to color the material poses no difference to the 300 GHz radar, in contrast to optical sensors. The 122 GHz radar appears to behave similarly to the 300 GHz device as far as plastic-sorting capabilities. While the sensor is still indifferent to filament dye, and is still able to detect when it is facing different sides of the block, it can now also detect the back-surface peak of the lower-density PLA design. This proves to be an added advantage, since it means the radar is able to penetrate the plastic in its entirety, allowing similar PLA designs to be used as data labels in the technique described in Section 5.1. Future work can be done to collect a larger database of measurements and materials, in order to extract appropriate thresholds for an automatic classification algorithm to group scanned objects according to their spectral properties.

Figure F.6: Spectrum comparison at 300 GHz of the black and red PLA designs when illuminating (a) their solid face or (b) their open face.

Figure F.7: Spectrum comparison of the two PLA pieces considered, when illuminating them on a (a) solid side or an (b) open side (exposing the piece's infill pattern), for **Configuration 3** of the 122 GHz radar.

Figure F.8: Spectrum comparison at 300 GHz of the HDPE and (a) black and (b) red PLA designs.

PLA filamento: ficha técnica

PLA Datasheet

PROPIEDAD	VALOR	UNIDADES	MÉTODO DE TESTEO
Propiedades generales			
Densidad	1,24	g/cm ³	ASTM D792
Propiedades mecánicas			
Módulo elástico en flexión	3600	MPa	ISO 178
Resistencia a la flexión	108	MPa	ISO 178
Dureza Sh D	85	Sh D	ASTM D2240
Propiedades térmicas			
Temperatura de flexión bajo carga	56	°C	ISO 75/2B
Temperatura de fusión	145-160	°C	ASTM D3418
Temperatura de transición vítrea	56-64	°C	ASTM D3418

Además de las propiedades descritas, hemos realizado los ensayos de tracción tanto de probetas inyectadas como de probetas impresas con nuestro PLA, para conocer las propiedades mecánicas de la pieza impresa final. En la siguiente tabla se recogen estos resultados:

Propiedades mecánicas	Probetas inyectadas	Probetas impresas ^a	Probetas impresas ^b	Unidades	Método de testeo
Resistencia a la rotura en tracción	52	50	39	MPa	ISO 527
Deformación a rotura en tracción	5	9	4	%	ISO 527
Módulo elástico en tracción	1320	1230	1120	MPa	ISO 527

^a Estirado en dirección paralela a las capas.
^b Estirado en dirección perpendicular a las capas

Probetas	Altura de capa (mm)	Grosor pared exterior (mm)	Grosor parte inferior/superior (mm)	Densidad del relleno (%)	Diseño de impresión
Probeta impresa en dirección paralela a las capas	0,2	2	0,2	100	
Probeta impresa en dirección perpendicular a las capas	0,3	1	0,3	100	

F.2 Analysis of Absorption in Common Fabrics

The analysis of material absorptions in the mm-wave range, initially applied to plastics in Section F.1 has been extended to everyday fabrics and textiles, using the same measurement set-up. The motive behind this extension is to study the effect of clothing placed on targets scanned by the imaging radars. In addition to blurring the resolution of the image (due to wrinkles, bagginess, and other factors), the addition of clothing layers creates increased losses from absorption. This can lower the received power from hidden threats located beneath these layers, even causing them to vanish, if the attenuation is large enough. Nonetheless, by previously measuring the standalone absorption effects of these materials, their effect on the image could be compensated in the processing stages. This objective has also been pursued by other research groups throughout the literature [457, 458].

For this analysis several pieces of clothing, with different thicknesses and made from various materials, have been used. The selection is limited in number to reduce the extension of the measurements, but most common textiles have been taken into account. A complete list of the clothings is presented in Table F.1.

Table F.1: Clothing pieces inspected with the 300 GHz and 122 GHz radars

Clothing type	Material	Thickness (mm)	Reference
Short-sleeve T-shirt	100% polyester	0.5	<i>Running T-shirt</i>
Jean pants	100% cotton denim	2	<i>Jeans</i>
Long-sleeve T-Shirt	100% cotton	0.5	<i>Black T-shirt</i>
Sweater	70% cotton + 30% nylon	1.5	<i>Green sweater</i>
Sweater	80% cotton + 20% acrylic	1.5	<i>Blue sweater</i>
Long-sleeve T-shirt	95% cotton + 5% spandex	1	<i>Brown T-shirt</i>

To measure the attenuation due to clothing a metallic sheet at 32.5 cm is once again used as reference target, and the different clothing pieces are suspended halfway between the sheet and the last focusing lens using a plastic 3-D printed frame, as illustrated by the diagram in Fig. F.9. At 300 GHz, the frame is centered with the horn antenna and lenses, to ensure that the radar illuminated only the textile and not the plastic. For the 122 GHz measurements, the set-up used in Section 5.1 is repeated, with the addition of the frame holder for the textiles. As was the case for the plastic inspections of Section F.1, the processing stages for all measurements in this analysis are those detailed in Section 5.1.1.

Figure F.9: Diagram of the measurement set-up for the textile absorption analysis with the 300 GHz radar, using the focusing scheme from Fig. 5.4 and placing a plastic frame between the lenses and target to hold the textiles.

A comparison among the 300 GHz spectrums from the metallic sheet and all textiles considered, may be seen in Fig. F.10. The first result that stands out is that most of the fabrics (except the *Blue sweater*, with just 0.05 dB losses) register higher returned power than the metallic sheet. The *Green sweater*, *Brown T-shirt*, *Black T-shirt*, *Jeans*, and *Running*

T-shirt present 0.37 dB, 0.55 dB, 0.61 dB, 1.39 dB, and 2.92 dB increases, respectively. This effect is entirely due to the joint aftermath of specular reflection from the sheet and diffuse reflection from the clothing. The smooth surface of the metal can redirect reflections slightly away from the radar if not perfectly aligned with the lenses, but when that redirected beam reaches the rough fabric of the clothing it spreads out, allowing the radar to capture more radiation. This sequence is depicted in Fig. F.11.

Figure F.10: Spectrum comparison at 300 GHz of a reference metallic sheet and all clothing pieces inspected.

Figure F.11: Diagram comparison of (a) purely specular reflection from a slightly-tilted metallic sheet with no clothing in front, and (b) the diffuse scattering created when placing the clothing.

To test out this theory, the measurements in Fig. F.10 are repeated without the use of the lenses. Without the lenses, the beamwidth widens and thus, for small tilts, the radar is still capable of capturing redirected reflections from the metal. The lens-less results are shown in Fig. F.12. In this case, the metallic sheet now presents the most powerful reflection (equal to the *Black T-shirt* scenario) as foreseen. The strongest attenuation (1.53 dB) comes from the the denim material of the *Jeans*, which is to be expected due to its larger thickness and absorption at 300 GHz (as measured in [23]). The *Green sweater*, *Blue sweater*, and *Brown T-shirt* all present similar attenuations (0.97 dB, 0.33 dB, and 0.47 dB, respectively) due to their comparable cotton content. Meanwhile, the *Running T-shirt*, although made completely from polyester (which has greater absorption than cotton for similar densities [23]), has a small thickness, creating only 0.56 dB in attenuation. The small thickness of the *Black T-shirt* is also the cause of its null absorption registered.

The spectral comparison among a reference metallic sheet (placed at 50 cm) and all fabrics considered (placed at 30 cm), for the 122 GHz radar is shown in Fig. F.13. At first glance, it

Figure F.12: Spectrum comparison at 300 GHz of a reference metallic sheet and all clothing pieces inspected, without using the focusing lenses.

can be noted that the specular reflection problems encountered at 300 GHz are not present, due to the removal of the focusing lenses when using the MMIC-based radar. All clothing scenarios present reflections less powerful than the metallic sheet, as expected when placing an attenuating layer in front of the target. According to Fig. F.13, transmissivity losses seem to be closely associated to clothing thickness. The *Running T-shirt* and *Black T-shirt* (both the narrowest fabrics, at 0.5 mm width) present very similar and low attenuations (0.24 dB and 0.71 dB, respectively). Following this trend, the next lowest attenuation (0.78 dB) is registered by precisely the third thinnest clothing piece (*Brown T-shirt*, with a 1 mm width). Larger losses (2.25 dB, 2.25 dB, and 2.40 dB) are created by the last three wider fabrics (*Jeans*, *Green sweater*, and *Blue sweater*, respectively). A summary of the transmissivity losses created by the fabrics for all frequency band considered is presented in Table F.2. Globally, the 122 GHz seems to present larger losses than the 300 GHz range, which goes against the common understanding that transmissivity increases with wavelength.

Figure F.13: Spectrum comparison of a reference metallic sheet and all fabrics considered, for) **Configuration 3** of the 122 GHz radar.

In conclusion, the low attenuations registered at 300 GHz would seemingly indicate that

little effect is caused on imaged targets by a clothing layer. Nonetheless, from the measurements it can also be concluded that textiles produce effects beyond simple attenuation. Depending on the focusing mechanism used, the introduction of fabrics over highly-specular threats (such as any metallic weapon) can cause great variations in the reflections received. This can be beneficial in terms of increasing received power, but can also be detrimental if the added scattering deforms the threat's shape in the image. Additionally, the clothing pieces used in these measurements have been stretched out to avoid uncontrolled wrinkles and folds, which are on the other hand extremely common when actually worn (producing additional losses and image distortion). The 122 GHz sensor receives a different attenuation ranking than the 300 GHz radar, with an apparent stronger dependence on clothing thickness. Contrary to the expected increase in transmissivity, the losses generated at the lower frequency band are higher than at the 300 GHz range. From an imaging point of view, this means that scans with a 122 GHz imager built around the MMIC-based radar will suffer more from clothing layers added to the target. As with the plastic study, future work should be focused on collecting a larger database of measurements and materials, in order to extract a ranking of fabrics according to their transmissivity losses. In addition, more textile variables (such as density and pigmentation) can be taken into account, and evaluations with even lower frequencies should be performed to test whether transmissivity losses continue to increase.

Table F.2: Transmissivity losses for all fabrics considered at 300 GHz and 122 GHz

Clothing reference*	300 GHz losses (dB)	122 GHz [†] losses (dB)
<i>Running T-shirt</i>	0.56	0.24
<i>Jeans</i>	1.53	2.25
<i>Black T-shirt</i>	0	0.71
<i>Green sweater</i>	0.97	2.25
<i>Blue sweater</i>	0.33	2.40
<i>Brown T-shirt</i>	0.47	0.78

*According to Table F.1.

[†]In **Configuration 3**.

F.3 Future Work

The analyses of Sections F.1 and F.2 allow for very diverse upgrades and future research. In this regard, further improvements in the set-up (such as evolving towards 3-D imaging, as mentioned in Section 5.1) and depth of the information extracted are the foremost needs.

For material evaluation and testing the next step in future work should be focused on the development of electromagnetic parameter-extraction algorithms. The extraction of a material's permittivity and permeability can allow its characterization, which can be useful in multiple NDT applications where an item's composition is questioned. Extraction techniques can generally be based on two approaches: resonant or non-resonant. Resonant methods extract permittivity and permeability from the resonant frequency of the material sample, while non-resonant approaches extract the parameters from the changing impedance that the transmission and reflection waves encounter after the sample is radiated. In the case of the radar described in this thesis, the extraction would have to be through non-resonant methods.

Within the non-resonant approaches several set-ups are commonly used, such as the examples shown in Fig. F.14. One of these set-ups is the transmission line approach, in which the sample is placed in a transmission line in the path of the travelling waves. A variant of this technique, directly achievable with the developments of this thesis, uses free-space as the transmission line. This method is usually executed with the use of a Vector

Network Analyzers (VNA), through the measurement of the S-parameters of a material sample placed between the two ports (where horn antennas are connected). Nonetheless, various examples of 1-port testing can be found, based on extraction of the S_{11} parameter only [463,464]. A disadvantage of reflection-only testing is that to accurately extract both permittivity and permeability two independent measurements of the sample are needed. To ensure independence, the mathematical expressions describing the measurements cannot be derived from one another. Several methods to force this condition can be used, such as using different sample thicknesses or varying the position or distance of a reference plane

Figure F.14: Examples of non-resonant-based set-ups for permittivity and permeability extraction from a material's samples, where a VNA is used as the radiating source. (a) is the open-ended coaxial probe technique, (b) is a 2-port transmission and reflection method, and (c) is a 1-port reflection-only approach.

Annex G

Bandwidth and Size Improvements for the 300 GHz Imaging Radar

During the development of this thesis, several possible improvements to the maximum transmitted bandwidth for the 300 GHz imaging radar have been detected, which is originally limited to 27 GHz. These improvements demand certain modifications to the RF architecture, which have nonetheless been restricted to preserve the radar's homodyne architecture and individual components (whenever possible). The proposed improvements have been analyzed observing only frequency ranges; for a more detailed analysis the power levels needed (and the ripples caused) at each stage should also be calculated and simulated.

To describe the improvement opportunities, a diagram of the radar's full architecture is found in Fig. G.1, where each step in the architecture is numbered. The frequencies currently used, and the maximum available, for each numbered step can then be consulted in Table G.1.

As can be seen from the table, none of the different steps throughout the architecture are limited to the current 27 GHz bandwidth, which means that an increase is possible. Because the TX and RX chains use different components, the possible increments in bandwidth depend on if the radar is used in its original or compact configuration [123].

If the radar is used in its compact setting, only the RX mm-wave chain is necessary. In this configuration, an increase to 55 GHz of bandwidth is possible. The limiting stages are as follows:

- The most bandwidth-restricting component is the active filter in the IF path, limiting to 32.85 GHz of bandwidth. The active filter is composed of a custom-made microstrip bandpass filter (from 0.8 - 14.6 MHz) and a commercial amplifier¹ (from 0 - 6 GHz) [127]. To improve the maximum bandwidth in this stage, the filter must simply be redesigned to increase its pass band.
- The next limiting component is the ADC, which restricts the RF bandwidth to 45 GHz because of its 40 MS/s maximum sampling rate. A simple solution to this obstacle, aside from a new faster ADC, is the use of bandpass sampling around the target's range (as in Section C).

¹Model SGA-3486Z from Qorvo.

Figure G. 1: Block diagram of the 300 GHz imaging radar with numbered bandwidth steps (in red), and the bandwidths and frequencies currently used in the system (in light blue).

Table G.1: Numbered steps in Fig. G.1 with currently used and maximum bandwidths (in GHz)

Number	Component	Input limit	Output limit	Used input	Used output	Max. BW*
1	DDS [432]	0.5 - 3.5 [†]	0 - 1.75	2.4	0.5 - 1	81
2	DDS output filter [432]	0 - 2.7	0 - 2.7	0.5 - 1	0.5 - 1	145.8
3	Equalizer [121]	0 - 1.2	0 - 1.2	0.5 - 1	0.5 - 1	64.8
4	Mixer [121]	0 - 2.5	4 - 7	0.5 - 1	5.3 - 5.8	162
5	Amplifier [121]	0 - 6	0 - 6	5.3 - 5.8	5.3 - 5.8	324
6	Bandpass filter [121]	5.1 - 5.8	5.1 - 5.8	5.3 - 5.8	5.3 - 5.8	37.8
7	x3 Multiplier [121]	3 - 5.8	9 - 17.4	5.3 - 5.8	15.9 - 17.4	151.2
8	Bandpass filter [121]	15 - 17.5	15 - 17.5	15.9 - 17.4	15.9 - 17.4	45
9	Coupler [121]	4 - 18	4 - 18	15.9 - 17.4	15.9 - 17.4	252
10	x6 Multiplier (Annex B.1)	12.5 - 18.33	75 - 110	15.9 - 17.4	95.4 - 104.4	105
11	Isolator [121]	75 - 110	75 - 110	95.4 - 104.4	95.4 - 104.4	105
12	MPA [121]	90 - 110	90 - 110	95.4 - 104.4	95.4 - 104.4	60
13	x3 Multiplier [121]	95 - 105	285 - 315	95.4 - 104.4	286.2 - 313.2	30
14	Horn antenna [121]	220 - 325	220 - 325	286.2 - 313.2	286.2 - 313.2	105
15	Beamsplitter [121]	100 - 400	100 - 400	286.2 - 313.2	286.2 - 313.2	300
16	SHM [121]	275 - 330 [‡]	0 - 8*	286.2 - 313.2	0 - 8	55
17	Isolator [121]	110 - 170	110 - 170	143.1 - 156.6	143.1 - 156.6	120
18	x3 Multiplier [121]	136.67 - 56.67	110 - 170	47.7 - 52.2	143.1 - 156.6	120
19	Attenuator [121]	40 - 60	40 - 60	47.7 - 52.2	47.7 - 52.2	120
20	MPA [121]	45 - 57	45 - 57	47.7 - 52.2	47.7 - 52.2	72
21	x3 Multiplier [121]	13.3 - 20	40 - 60	15.9 - 17.4	47.7 - 52.2	120
22	Active filter [127]	(0.8 - 14.6) · 10 ⁻³	(0.8 - 14.6) · 10 ⁻³	(0.8 - 14.6) · 10 ⁻³	(0.8 - 14.6) · 10 ⁻³	32.85 [⊗]
23	ADC	(0 - 20) · 10 ⁻³	(0 - 20) · 10 ⁻³	(0 - 20) · 10 ⁻³	(0 - 20) · 10 ⁻³	45 [⊗]

*Equivalent bandwidth at RF range, not at the component's frequency range.

[†]Input clock frequency.[‡]RF port.

*IF port.

⊗With a 150 μ s chirp and a beat frequency at 10 m, according to (1.18).

- The following limiting component is the SHM, which restricts bandwidth in its RF port to 55 GHz, at 275 - 330 GHz. These limits and bandwidth are perfectly compatible with the remaining components in the mm-wave chain. Therefore, no modifications are necessary to the RX chain.
- To use the full 275 - 330 GHz RF bandwidth, the microwave subsystem's output must be modified from 15.9 - 17.4 GHz to 15.27 - 18.33 GHz. Several of the custom-made filters would need to be redesigned for this bandwidth, but most importantly, the current x3 multiplier² cannot be used as it would limit its output to 15.27 - 17.4 GHz (38.34 GHz of RF bandwidth). Because no other commercial tripler for the desired band was found, a new microwave architecture was proposed to generate the desired bandwidth [432]. A block diagram of the architecture is shown in Fig. G.2. This architecture is based around the use of the DDS's image bands outside of the first Nyquist zone, and commercial doublers instead of triplers. In this new architecture the lowpass filter after the DDS is changed to a bandpass filter to select the desired image band. The filter would also need to be designed to equalize the higher power ripple the chirp experiences in the image bands [432]. After the filter, two sets of amplifiers and doublers³ shift the chirp to the desired band.

Figure G.2: Proposed new microwave subsystem for the 300 GHz imaging radar using its DDS's image bands.

For the radar's original configuration, in which both the RX and the TX mm-wave chains are used, several bandwidth improvement options are proposed:

1. The first proposal is an immediate increase to a 28.26 GHz bandwidth (at 284.94 - 313.2 GHz), without needing to modify any component in the architecture. For this improvement the limiting component is once again the tripler in the microwave subsystem.
2. The second proposal is an increase to a 30 GHz bandwidth (at 285 - 315 GHz), through modification of the microwave subsystem. This subsystem would need to be modified with the same proposed architecture as in Fig. G.2, to avoid the use of the limiting tripler. The subsystem's new output would be at 15.833 - 17.5 GHz. In this proposal, after modifying the microwave architecture, the limiting component is the tripler in the TX mm-wave chain.
3. The final proposal is an increase to 60 GHz of bandwidth, but it requires changes in the mm-wave chains, which can be expensive. The proposed increase also requires changing the microwave subsystem to the architecture seen in Fig. G.2. After modifying the microwave subsystem the mm-wave changes would be as follows:

²Model HMC916 from Hittite.

³For example, model KC2-50+ from Mini-Circuits for the first doubler and model HMC814LC3B from Analog Devices for the second doubler.

- The most limiting component in the mm-wave chains is the TX chain's tripler, which can be substituted for another COTS tripler from the same manufacturer series⁴ which allows a 175 GHz bandwidth at 325 - 500 GHz. This change would impose a much higher central frequency for the radar's transmission, which seemingly goes against the current trend in imaging radars to work at the 300 GHz window. The benefit of working at a higher frequency is better cross-range resolution for a fixed antenna size [according to (3.3)], at the cost of worse atmospheric attenuation and clothing transmissivity.
- Because of the higher RF range, the horn antenna would also need to be changed from the current one (which can be used only at 220 - 325 GHz) to another COTS alternative⁵ which works at 325 - 500 GHz.
- In the RF path, the beamsplitter currently used would then limit the bandwidth to 75 GHz at 325 - 400 GHz.
- The next most-restricting component would then be the MPA in the TX chain, which would limit the RF bandwidth to just 5 GHz around 325 - 330 GHz. Because no COTS alternatives to this MPA at the desired band have been found, the proposed solution is to use a COTS x12 multiplier chain⁶ to substitute the TX chain up to the tripler. This device accepts an input from 9.17 - 14.16 GHz and multiplies it by a x12 factor. With this change, the RF bandwidth would then be 69.88 GHz at 330.12 - 400 GHz.
- After the previous modifications, the next limiting component is the SHM, which is placed in the RX chain. This device only accepts 275 - 330 GHz in its RF port, so it is not compatible with the proposed 330.12 - 400 GHz bandwidth. Because no other COTS SHM at the desired band could be found, a two-stage conversion using a harmonic mixer⁷ and a SHM⁸ is proposed (as seen in Fig. G.3). With this new conversion architecture, the proposed 69.88 GHz bandwidth can be used, if the LO frequency is set between 70 - 110.12 GHz.
- For the new SHM's output, which will be between 220 - 330 GHz, the RX chain's isolator and tripler pose no problem. The chain's attenuator will limit the SHM output range to 240 - 330 GHz, but the proposed 69.88 GHz bandwidth can still be used if the LO frequency for the harmonic mixer is now between 70 - 90.12 GHz. On the other hand, the RX chain's MPA further limits the SHM output range to 270 - 330 GHz, so the RF bandwidth is now limited to 60 GHz, between 330.12 - 400 GHz, depending on the exact value of the LO frequency during conversion.

As a summary, Table G.2 presents all the possible improvements.

In addition to the detected improvements in bandwidth, the radar's RX mm-wave chain can also be reduced in size. This can be done via an update on the technology used for the components up to 52 GHz, switching from the current bulky waveguide devices to commercial MMICs [127]. The use of MMICs up to 52 GHz would allow the removal of the RX chain's first tripler, its MPA, and its attenuator (leaving only the isolator, SHM, and horn antenna). Furthermore, the substitution of these three components for MMICs is also a large reduction in cost. The proposed architecture uses two types of MMICs; a saturated amplifier working as tripler⁹ and a high-power amplifier¹⁰ to amplify its output.

⁴Model FM3 325-500 from Radiometer Physics GmbH.

⁵For, example, model FH-PP-500 from Radiometer Physics GmbH (the same manufacturer as the current antenna).

⁶Model WR6.5AMC-I from Virginia Diodes.

⁷Model HM3 325 - 500 from Radiometer Physics GmbH, which accepts a 325 - 500 GHz RF input.

⁸Model SHM 220 - 330, from Radiometer Physics GmbH, which accepts a 220 - 330 GHz RF input.

⁹Model AMMC-5040 from Avago.

¹⁰Model HMC1144 from Hittite.

Figure G.3: Proposed new RF-to-IF conversion for the 300 GHz imaging radar using two mixers.

Table G.2: Summary of the proposed bandwidth improvements for the 300 GHz radar.

Configuration	BW (GHz)	BW Increment	Micro. changes? [†]	Milli. changes? [‡]
Compact	55	103.70%	Yes	No
Original	28.26	4.67%	No	No
Original	30	11.11%	Yes	No
Original	60	122.22%	Yes	Yes

[†]Changes to microwave subsystem necessary.

[‡]Changes to the mm-wave chains necessary.

Furthermore, custom-made frequency multipliers at 100 GHz and 300 GHz are currently being developed by members of the GMR to provide wider bandwidths, as well as possible improvements to transmitted power and device size [96].

Acronyms

ADC Analog-to-Digital Converter
ATD Automatic Threat Detection
ATR Automatic Threat Recognition
AWGN Additive White Gaussian Noise
BEGRS Bifocal Ellipsoidal Gregorian Reflector System
BRDF Bidirectional Reflectance Distribution Function
CNN Convolutional Neural Network
COTS Commercial Off-The-Shelf Components
CUDA Compute Unified Device Architecture
CW Continuous Wave
DANL Displayed Average Noise Level
DBSCAN Density-Based Spatial Clustering of Applications with Noise
DDS Direct Digital Synthesizer
DHS Department of Homeland Security
DRO Dielectric Resonant Oscillator
DSP Digital Signal Processor
DWT Discrete Wavelet Transform
ECG Electrocardiogram
EPS Expanded Polystyrene
FFT Fast Fourier Transform
FMCW Frequency-modulated Continuous Wave
FoV Field of View
FPA Focal Plane Array
FPGA Field-Programmable Gate Array
GMR *Grupo de Microondas y Radar*
GPGPU General Purpose computing on Graphics Processing Unit
GPR Ground-Penetrating Radar
GPS Global Positioning System
GPU Graphics Processing Unit
HD-AIT High Definition-Advanced Imaging Technology
HDPE High-Density Polyethylene
HPBW Half-Power Beamwidth
IC Integrated Circuit
ICP Iterative Closest Point
IDR Instantaneous Dynamic Range
IED Improvised Explosive Device
IF Intermediate Frequency
IFFT Inverse FFT
I/O Input/Output
IR Infrared
IRQ Interrupt Request
ISM Industrial, Scientific, and Medical
JPL Jet Propulsion Laboratory
LFM Linear Frequency Modulated
LIDAR Light Detection And Ranging
LNA Low-Noise Amplifier
LO Local Oscillator
MIMO Multiple-Input and Multiple-Output
MMIC Monolithic Microwave Integrated Circuit
MPA Medium-Power Amplifier
NDT Non-Destructive Testing
NFC Near-field Communication
OUT Object Under Test

PCB Printed Circuit Board
PCH Platform Controller Hub
PFD Phase-Frequency Detector
PLA Polyactic Acid
PLL Phase-Locked Loop
PLO Phase-Locked Oscillator
PMCW Phase-Modulated Continuous-Wave
PO Physical Optics
PPG Photoplethysmogram
QR Quick Response
RCS Radar Cross-Section
ReLU Rectified Linear Unit
RF Radiofrequency
RFID Radio Frequency IDentification
RX Receiver
SAR Synthetic Aperture Radar
SHM Subharmonic Mixer
SIMT Single Instruction Multiple Thread
SLI Scalable Link Interface
SNR Signal-to-Noise Ratio
TCP Transmission Control Protocol
TRX Transceiver
TX Transmitter
TSA Transportation Security Administration
UDP User Datagram Protocol
UPM *Universidad Politécnica de Madrid*
VCO Voltage-Controlled Oscillator
VNA Vector Network Analyzer
VTK Visualization Toolkit

Firmado: Federico Antonio García Rial.

Ingeniero de Telecomunicación.

Madrid, a de de 2019