

RESEARCH ARTICLE

# Techno-economic analysis for smart hangar inspection operations through sensing and localisation at scale

A. Plastropoulos<sup>1</sup>, N. P. Avdelidis<sup>2</sup> and A. Zolotas<sup>1</sup>

<sup>1</sup>Faculty of Engineering and Applied Sciences, Centre for Assured and Connected Autonomy, Cranfield University, Cranfield, UK

<sup>2</sup>Department of Aeronautics and Astronautics, School of Engineering, University of Southampton, Southampton, UK

**Corresponding author:** A. Plastropoulos; Email: [a.plastropoulos@cranfield.ac.uk](mailto:a.plastropoulos@cranfield.ac.uk)

**Received:** 20 September 2025; **Revised:** 21 November 2025; **Accepted:** 25 November 2025

**Keywords:** aircraft maintenance; hangar of the future; localisation; monitoring; MRO; smart infrastructure

## Abstract

The accuracy, robustness and affordability of localisation are fundamental to autonomous robotic inspection within aircraft maintenance, repair and overhaul (MRO) hangars. Hangars typically have high ceilings and are predominantly steel-framed structures with metal cladding. Because of this, they are regarded as GPS-denied environments, characterised by significant multipath effects and strict operational constraints, which together form a unique challenging setting. The lack of comparative techno-economic benchmarks for localisation technologies in such environments remains a critical gap. Addressing this, the paper presents the first techno-economic analysis that benchmarks motion capture (MoCap), ultra-wideband (UWB) and a ceiling-mounted camera (CMC) system across three operational scenarios: robot localisation, asset monitoring and surface defect detection within a single-bay hangar. A two-stage optimisation framework for camera selection and placement is introduced, which couples market-based camera-lens selection with an optimisation solver, producing camera layouts that minimise hardware while meeting accuracy and coverage targets. The consolidated blueprints provide quantification of the required equipment and its performance: 15 global-shutter GigE cameras are adequate for drone localisation, 9 cameras meet the requirements for on-bay monitoring and 49 high-resolution cameras facilitate defect mapping of the upper airframe surfaces for midsize defects. Across these scenarios, the study reports indicative performance and cost envelopes: a MoCap installation delivers submillimeter localisation at an estimated £190k per bay, UWB delivers centimetre-level tracking for around £49k and the proposed CMC system layouts achieve task-specific coverage with costs in the £9k–£77k range. The analysis equips MRO planners with an actionable method to balance accuracy, coverage and budget, demonstrating that an optimised CMC system can deliver robust and cost-effective sensing for next-generation smart hangars.

## Nomenclature

MRO	maintenance repair and overhaul
UWB	ultra-wideband
MoCap	motion-capture
DOF	degrees of freedom
PoE	Power over Ethernet
EKF	extended Kalman filter
ToA	time-of-arrival
TDoA	time-difference-of-arrival
IMU	inertial measurement unit
FOD	foreign object debris
DXF	Drawing eXchange Format
OEM	original equipment manufacturer

NIC	network interface cards
MTU	maximum transmission unit
LpBinary	linear-programming of binary type
CBC MILP	COIN-OR branch-and-cut – mixed-integer linear programming
SLAM	simultaneous localisation and mapping
MILP	mixed-integer linear programming
ROS	robot operating system
GigE	gigabit ethernet
LiDAR	light detection and ranging
USB	universal serial bus
CCTV	closed-circuit television
SVG	Scalable Vector Graphics
SPF	small form-factor pluggable
UTP	unshielded twisted pair
Cat6	category 6
TOF	time of flight
RTT	round-trip time
NDT	nondestructive testing
GSD	ground sampling distance
FoV	field of view
TRL	technology readiness level
CMC	ceiling-mounted camera
RF	radio frequency

## 1. Introduction

Accurate localisation is fundamental to enabling autonomous robotic inspection in MRO hangars, where metallic structures, extensive multipath effects and strict operational constraints define a unique and challenging environment [1]. Existing localisation technologies, including infrared MoCap, UWB, and camera-based systems or on-board simultaneous localisation and mapping (SLAM) approaches, offer different trade-offs in terms of achievable accuracy, infrastructure complexity, cost and robustness to occlusion and interference [2]. Sensor fusion frameworks that combine vision, inertial and UWB data can improve robustness and deliver cost–accuracy trade-offs, as demonstrated in large-scale warehouse deployments. However, their performance and economic viability remain highly dependent on environment-specific factors [3]. In the context of this research, monitoring is defined as the real-time tracking of assets such as ground support equipment, tool trolleys and personnel within the bay. Localisation focusses on determining the real-time position and orientation of ground and aerial platforms operating in proximity to the aircraft. Finally, artefact detection is addressed, with an exploration of how sensing can facilitate the identification of surface defects and other critical features on the airframe. The framing of these macroscopic and microscopic needs together sets the stage for the optimisation framework, comparative experiments and cost analyses developed in the remainder of the paper.

Despite technical advances, a notable lack of comprehensive, real-world techno-economic analyses remains, specifically focused on aircraft hangar deployments. The available literature provides only partial benchmarking or component-level comparisons for individual or hybrid localisation modalities [4, 5], with little empirical evidence on their robustness to the full spectrum of hangar-specific challenges such as dynamic occlusion, specular reflections and integration with existing maintenance workflows. To date, no studies have presented a holistic side-by-side evaluation of MoCap, UWB and vision-based solutions in an operational aircraft hangar context. This persistent gap underscores an urgent need for domain-specific comparative studies, including rigorous assessments of cost, accuracy and integration, to inform the reliable and scalable deployment of robotic inspection systems in aviation environments. The economic impact of daily maintenance practices has begun to be quantified by recent hangar-focused

studies; for instance, significant rework costs in narrow-body bays can be avoided by improving technicians' awareness of composite repair, as demonstrated by Jong et al. [6]. Across the timeline, Moenck et al. [7] outline how the forthcoming Industry 5.0 automation could reshape the trade-offs of labour hours and logistics on large MRO campuses. At the same time, the classic aerodynamic analysis of the enclosed engine test hangars by Wallis and Ruglen still provides a valuable historical baseline for energy throughput economics [8].

In summary, this study offers five significant contributions: (i) it presents the inaugural techno-economic analysis that evaluates MoCap, UWB and CMC systems in parallel for full-scale aircraft hangars; (ii) it introduces a two-stage optimisation framework for camera selection and placement that combines market-driven camera-lens selection with a mixed-integer linear programming (MILP)-based set-cover placement, resulting in the minimal number of cameras needed; (iii) it supplies quantified design-to-cost case studies converting three typical MRO tasks into specific bills of materials and cost estimates; (iv) determines the optimal balance for defect-detection accuracy, illustrating how ceiling cameras and drone close-ups converge at various defect sizes; and (v) provides the first cost/accuracy comparison between camera localisation and commercial UWB/MoCap systems for a conventional  $40 \times 50$  m bay. Collectively, these contributions deliver an actionable and comprehensive methodology for MRO decision-makers to select, size and cost localisation and inspection systems within large hangars. The framework demonstrated in the case studies is applied to a rectangular single bay hangar measuring  $40 \times 50$  m for narrowbody aircraft. The optimisation and coverage mapping processes are fully parametric in terms of bay width, length and ceiling height.

Inspired by the aviation industry's shift towards Industry 5.0, which sees mobile robots and AI-based decision support systems taking on routine maintenance duties, the hangar should transition from a passive shelter to a dynamic sensing platform. The end goal is a ceiling infrastructure dense enough to localise robots, track assets and even surface defects in real-time, but lean enough to be economically retrofitted into legacy bays. Against this backdrop, the remainder of the paper is organised as follows. Section 2 reviews the state-of-the-art in MoCap, UWB and CMC systems, clarifying their respective accuracy, cost and integration trade-offs. Section 3 introduces a two-stage optimisation framework for camera selection and placement that first selects a market-ready camera–lens pair and then solves a set-cover problem to minimise hardware while meeting resolution targets. Section 4 translates those algorithms into three design-to-cost frameworks: robot localisation, asset tracking and defect detection, each with a detailed bill of materials for practitioner guidance. It also presents MoCap and UWB implementation options for benchmarking.

## 2. Digital Sensing Options for the Smart Hangar

The hangar is a space where multiple maintenance tasks from various technical teams are performed simultaneously. Effective coordination, spatial awareness and asset management are crucial to increase productivity and minimise mistakes. In addition, precise localisation is crucial for evaluating applications involving robotics. Therefore, it is no surprise that when Airbus first introduced the concept of the Hangar of the Future, the localisation system was part of the suggested enabling technologies [9]. For this concept, the UWB system was chosen. The choice was reasonable since it provides a balance between cost and accuracy. Ubisense has publicly shared two related experiments. In the first, the system provides localisation data for a drone that performs aircraft inspection [10]. In the second study [11], the focus was on asset management and digital twins. An alternative method, which is not widely adopted in hangars, involves utilising a MoCap system. Although this provides greater accuracy, it is more expensive and requires a line of sight with a minimum of two cameras [12]. Yet another technique, still in the research stage but showing significant potential, is to explore the usage of a camera-based system enhanced with deep learning models capable of executing various tasks [13, 14]. The subsequent section outlines each approach by highlighting its fundamental functions, characteristics and constraints, along with its standard architecture. For the camera-based system, three candidate architectures are illustrated on the basis of specific target scenarios.

In this section, background is grouped by sensing family (MoCap, UWB, CMC systems), followed by use-case (localisation, monitoring, defect detection), with dimensional assumptions kept adjacent to each use-case for easy cross-reference to Section 3 (methods) and Section 4 (blueprints).

### **2.1. Optical motion-capture systems**

MoCap systems are technologies that estimate the 3D positions and orientations (six degrees of freedom (DOF) poses) of tracked objects using camera-based tracking of reflective or active markers. These systems operate by triangulating marker positions from multiple synchronised camera views and are widely used in fields such as robotics, biomechanics and animation. Commercial systems such as Qualisys, Vicon and OptiTrack use high-speed infrared cameras configured around a workspace to track marker-equipped objects with submillimeter precision. For robotic localisation tasks, especially in indoor settings, MoCap systems provide highly accurate pose data that can be integrated into robotic frameworks, such as those based on the robot operating system (ROS), enabling real-time control, mapping or trajectory following.

The MoCap systems use multiple infrared cameras and reflective markers to determine 3D positions in real-time. Reflective markers are small spheres (often coated with retro-reflective material) placed on the object of interest. The reflective coating ensures that the markers appear as bright spots in the camera images whenever IR light is reflected from them. For robot localisation, typically, multiple markers are affixed in a fixed configuration on the robot, forming a rigid body. A typical MoCap setup (Fig. 1(a)) involves several hardware components, such as infrared cameras with built-in IR LED strobe rings, synchronisation modules, data acquisition units and a central processing computer. The architecture ensures that all cameras capture movement simultaneously and feed data to a hub that computes the position and orientation of the robot at high update rates. For accurate 3D reconstruction, all cameras must capture frames in sync. A synchronisation unit ensures that each camera exposure is synchronised in time. This hardware unit maintains a consistent frame rate across all cameras, ensuring that the position of a fast-moving object is recorded at the same instant from all viewpoints. Optical MoCap cameras typically connect to a data acquisition network that handles both data transfer and power supply. This is often achieved via Power over Ethernet (PoE). A PoE network switch serves as the central hub to which all cameras connect via Ethernet cables. This switch provides DC power to the cameras via the same cable that carries data, simplifying installation (no separate power cords are required for each camera). An alternative method involves implementing a network of Ethernet-powered cameras connected in a daisy-chain configuration. Finally, in the centre of the system is the processing hub, typically a host computing unit that runs MoCap software. This unit serves as the central hub where all camera data are collected for real-time processing and storage.

In robotic applications, particularly in large indoor environments such as industrial inspection areas or aircraft hangars, the deployment of MoCap systems faces a number of challenges. These include maintaining camera visibility across wide areas, synchronising data from distributed cameras, mitigating occlusion or misalignment of markers, and ensuring system accuracy despite environmental interference such as lighting or reflective surfaces.

Among the studies reviewed, commercial systems have been extensively benchmarked under controlled conditions, showing static accuracy of approximately 0.15 mm and dynamic errors of less than 2 mm [15], although these results stem from experiments not explicitly conducted in large-scale industrial settings. Studies such as Ref. (16) have proposed optimised camera placement algorithms to improve tracking performance under challenging configurations, which is relevant to scalability.

A notable finding from the literature is that, despite the high performance of commercial systems, few papers rigorously evaluate their use in truly large-scale or industrially challenging environments. Some research explores alternative low-cost multicamera systems combined with an extended Kalman filter (EKF) or portable near-infrared camera approaches [17, 18], although these systems are typically validated only in small or medium-sized settings. Studies such as Ref. (19) demonstrate a comparison of fiducial-based SLAM benchmarking methods in large industrial-like spaces but do not fully address

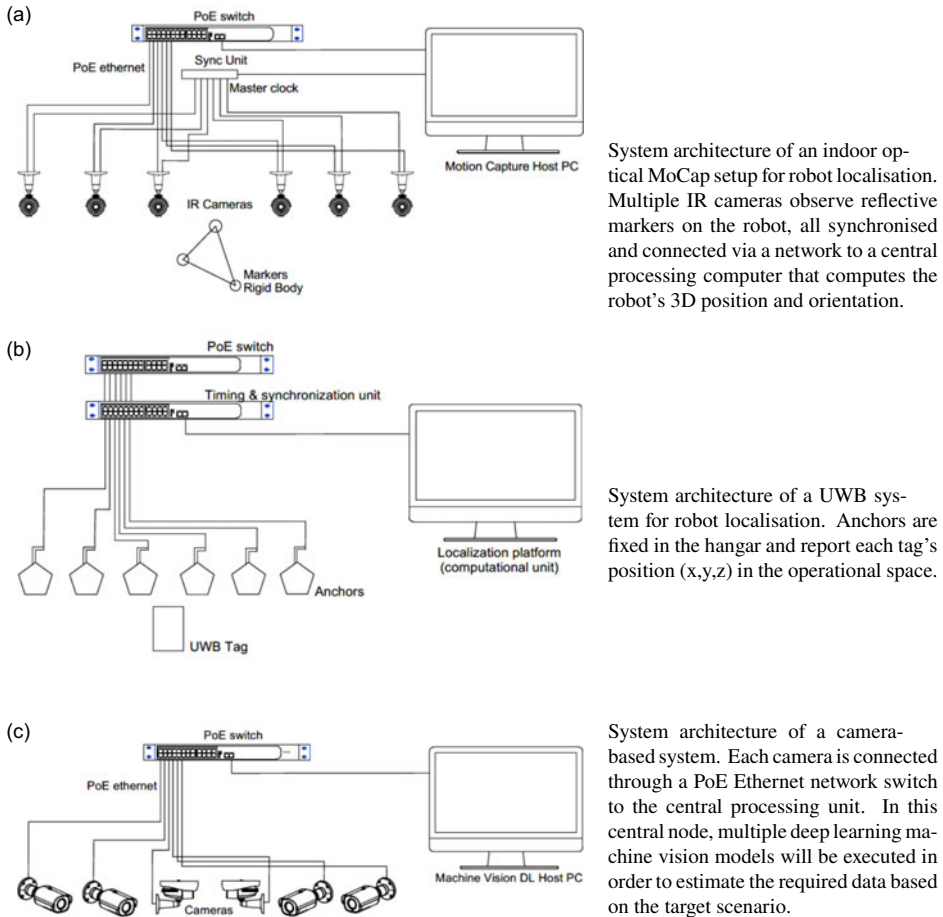


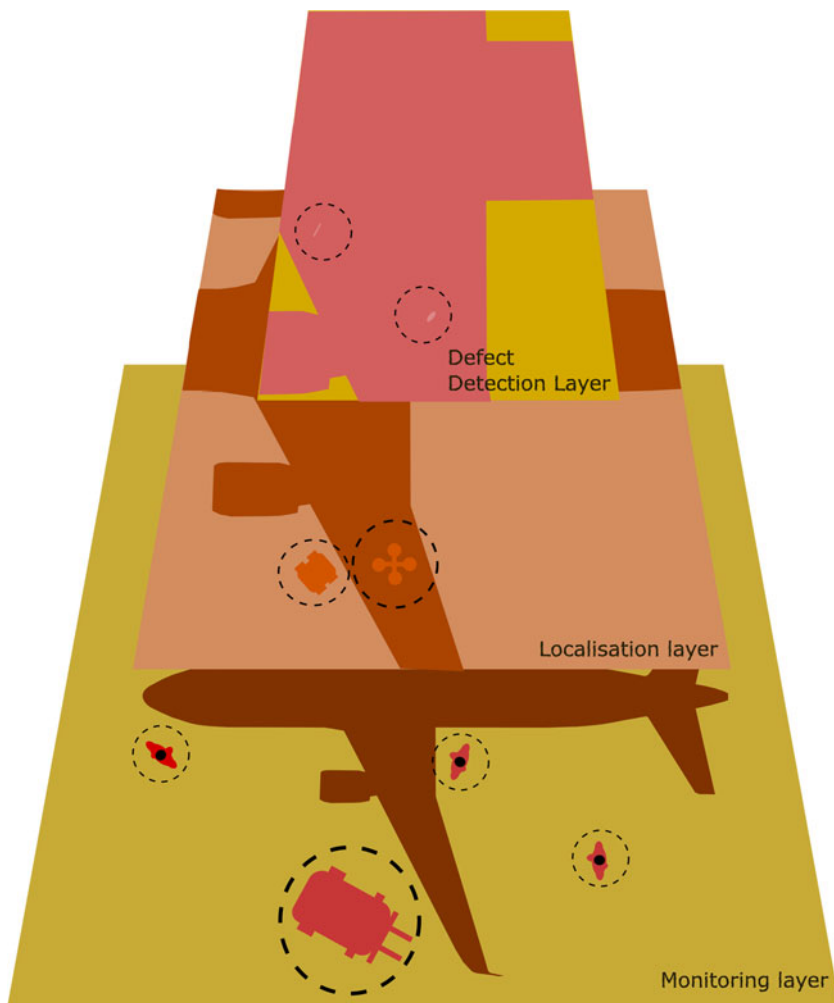
Figure 1. The three proposed system architectures.

scalability, calibration complexity or environmental robustness. In general, while commercial MoCap systems provide excellent precision, the literature reveals significant gaps in understanding how they perform and scale within expansive environments like aircraft hangars.

## 2.2. Ultra-wideband systems

UWB localisation systems have become efficient methods for achieving accurate indoor positioning within complex, RF-challenging environments, such as industrial sites. Generally, these systems employ time-of-arrival (ToA) or time-difference-of-arrival (TDoA) methods, which can deliver centimetre-level precision when conditions are optimal [20]. UWB systems operate by transmitting short-duration pulses over a wide frequency spectrum, enabling them to effectively mitigate multipath interference and signal fading, common challenges encountered by narrowband solutions in complex indoor environments.

Studies conducted in structured real-world environments, such as industrial plants [21], confirm the strong performance of UWB in environments with metallic surfaces and multipath interference. UWB systems offer key advantages over vision-based alternatives, including robustness to occlusions and lighting variations, as well as reduced computational complexity in localisation pipelines due to lightweight signal processing requirements [22]. However, the deployment of UWB systems also presents challenges, particularly for applications such as aircraft hangars. Multipath effects caused by reflective environments and dynamic obstructions (e.g., moving machinery or personnel) can degrade



**Figure 2.** The camera system could be configured to operate in three modes: monitoring, localisation, and defect detection.

precision if not properly addressed. Several approaches have been proposed to mitigate these issues, such as machine learning-based corrections [23] and multipath-assisted localisation techniques [24]. Scalability is a relative strength; systems have been successfully deployed in facilities exceeding 1,500 m<sup>2</sup> [25], but the complexity of the setup increases with the size of the system.

A typical UWB localisation system consists of multiple fixed UWB anchors installed at known positions within the environment and UWB tags mounted on the object of interest (Fig. 1(b)). The system estimates the position of the tag by measuring the ToF, TDoA or round-trip time (RTT) of UWB radio pulses exchanged between the tag and anchors. A localisation engine processes these timing measurements on an external computing unit to compute the tag's position through multilateration or filtering techniques such as Kalman or particle filters. Synchronisation between anchors is crucial for accurate timing. UWB systems are often complemented at the platform level using additional sensors such as an inertial measurement unit (IMU), light detection and range (LiDAR) or vision systems to improve localisation robustness and accuracy, particularly in non-line-of-sight or multipath conditions. For large-scale, high-precision applications like aircraft hangars, UWB presents a strong foundational technology, provided that deployment challenges are managed through intelligent design and calibration strategies.

**Table 1.** Summary of the modes of operation, the typical dimensions of the target objects and their velocity

Mode	Objects of interest	Typical L × W (m)	Typical H (m)	Velocity (m/s)
Localisation	Drones	0.5 × 0.5	4.5–6.5	0.5–1.5
Localisation	Wheeled platforms and quadrupeds	0.8 × 0.5	0.38–0.6	0.5–1.5
Monitoring	Scissor lifts, cherry pickers, forklifts, tugs	2.4 × 1.1	1.2–2.0	0.9–1.8
Monitoring	Scissor lifts, cherry pickers, forklifts, tugs, humans and trolleys	0.61 × 0.46	1.0–2.0	0.8–1.7
Defect detection	Small defects such as cracks and lightning strikes	0.01 × 0.01	Not applicable	Not applicable
Defect detection	Medium defects such as dents and peeled paint	0.04 × 0.04	Not applicable	Not applicable

### 2.3. Ceiling-mounted camera systems

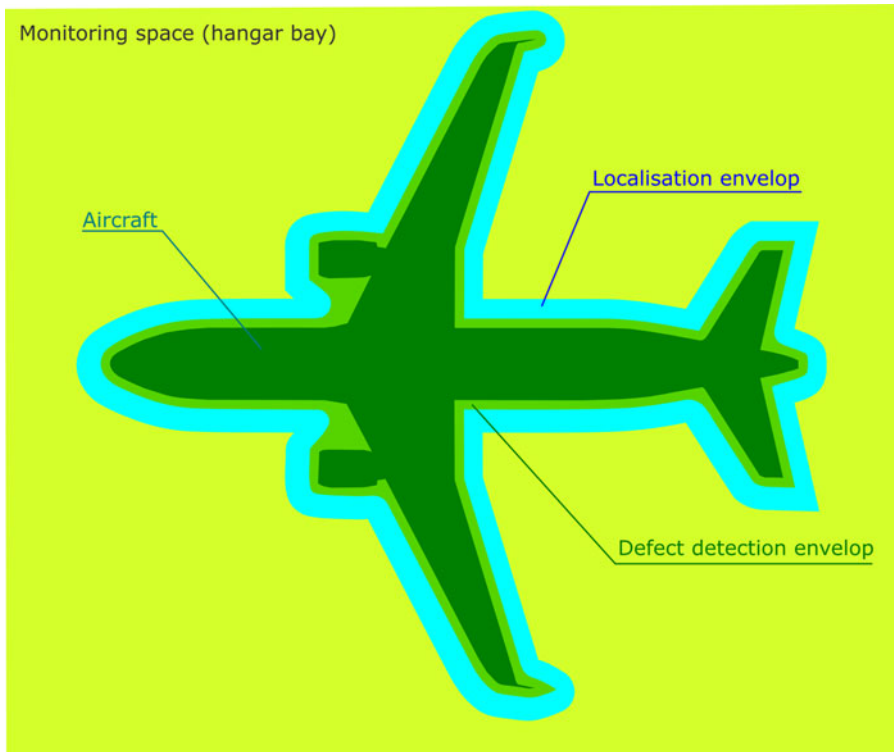
Utilising a camera-based setup provides versatility in detecting and tracking objects of interest. The system parameters, including sensor dimensions and the lens's field of view, can be modified to suit the desired application. The system is simpler compared to the other approaches, as it consists of cameras, PoE network switches, and the computing unit that executes the required algorithms (Fig. 1(c)). As the primary functionality of the system relies on the algorithms, this allows the system to be repurposed depending on the target scenario. There are three modes that can be considered using the camera system. As illustrated in Fig. 2, it can be used to locate robotic platforms, monitor assets or detect defects. In the following subsections, for each mode, the typical lateral dimensions of the target objects, the height and the velocity are defined. Every operational parameter is linked to the specifications of the camera. Table 1 summarises the assumptions that were analysed in the following subsections.

#### 2.3.1. Drone and ground platform localisation mode

Robotic platforms for inspection purposes can vary from aerial systems to ground units. The choice of platform is determined by the nondestructive-testing (NDT) sensors and the inspection target. Drones are well-suited for inspecting vertical and horizontal stabilisers, as well as the upper sections of the fuselage and wings. Meanwhile, landing gears and areas beneath the wings and fuselage are more suitable for ground-based robotic platforms. Depending on the demands and budget, options such as wheeled, tracked or even quadruped systems can be explored.

There are many approaches available for drones classification based on size, mass and altitude. In this study, the focus is on the frame size, which is often referred to as the diameter. The diameter of mini quad-copters usually ranges between 250 and 1,000 mm, with 500 mm being a common choice depending on the application [26]. An acceptable presumption for the drone's speed during inspection activities might fall between 0.5 and 1.5 m/s. Regarding wheeled robots, representative cases can be considered Husarion's Panther and Clearpath's Husky, with external dimensions (L × W × H) 810 × 850 × 370 mm and 990 × 698 × 372 mm, respectively [27, 28]. The most popular quadruped, which is the Boston Dynamics Spot robot, has external dimensions of 1,100 × 500 × 610 mm [29].

In the drone's localisation, the assumption is that the drone will fly 0.5–1.0 m above the aircraft surface, which translates to an altitude of 4.5–6.5 m. Regarding ground mobile platforms, it is acceptable to assume heights in the range between 380 and 600 mm. In terms of operational space, the focus is on covering the aircraft and the surrounding region. Figure 3 shows the aircraft along with the coverage envelope (blue indigo), which expands the boundary of the aircraft by 1 m. In conclusion, the focus in terms of distance from the ground is 4.0–7.0 mm for drones and 0.5–1 m for ground platforms.



**Figure 3.** Different candidate modes of operation for the camera-based system.

### 2.3.2. Asset and personnel tracking mode

Hangars are busy environments with a variety of ground support vehicles performing tasks during MRO operations. Some examples could be scissor lifts or cherry pickers ( $L \times W$ )  $2.4 \times 1.1$  m, forklifts  $3.8 \times 1.2$  m, ground power units  $3.3 \times 1.8$  m and tug tractors  $4.5 \times 2.0$  m. This list is not exhaustive, but it contains the most frequent vehicles found inside the hangar. On a different scale size, and not being a vehicle, tooling trolleys are also frequently around technicians performing maintenance with typical lateral dimensions of  $0.95 \times 0.46$  m. A significant factor to consider in the monitoring mode is the inclusion of humans. The analysis presented in Ref. (30) suggests that humans can be represented as ellipses, where the minor axis corresponds to body depth and the major axis corresponds to shoulder width. The shoulder width is estimated to be 610 mm, while the body depth is approximately 457 mm. Typical indoor speed limits for support vehicles moving inside the hangar range between 0.9 and 1.8 m/s, when humans are present and driving in free space, respectively [31, 32]. In humans, walking speed is an essential gait parameter that indicates the functional state and overall health of a person. This metric is typically assessed by the distance covered over a period of time, with speeds varying from a leisurely 0.82 m/s to a brisk 1.72 m/s [33]. In terms of height, the lower and upper surfaces of the support vehicles are within the range of 1.2 and 2.0 m (Fig. 4), and for tool trolleys around 1.0 m. If humans are included, according to Ref. (34), the average height of an adult male could be approximated to 1.7 m, which falls within the aforementioned range. The operational space is the area surrounding the aircraft where maintenance tasks are performed (highlighted in yellow in Fig. 3). According to building manufacturers [35], the standard size of a single-bay narrow-body maintenance space varies from  $40 \times 50$  m to  $60 \times 60$  m.

### 2.3.3. Surface-defect detection mode

Commercial aircraft surfaces are prone to a range of defects that can occur during manufacturing, use and maintenance. These defects, ranging from tiny corrosion pits to noticeable dents and cracks, can



**Figure 4.** The front view of relevant assets inside a hangar in real scale (aircraft Airbus A320).

affect structural integrity and safety. The most common types of defects observed on aircraft exterior surfaces can be broadly grouped as follows:

- **Impact-driven dents.** Among the most frequently encountered surface defects are dents arising from hail impacts, tooling mishaps, foreign object debris (FOD) or bird impacts. The dimensions of the dents depend on the energy of the impact, the properties of the material, and the environmental conditions. Hail damage produces dents with depths related to the diameter and impact angle of the hailstone [36]. Chen et al. [37] analysed damage records spanning more than a decade and reported typical lateral sizes of 38–50 mm; the same study notes that manual visual inspection usually identifies dents larger than 10 mm.
- **Coating and paint degradation.** The deterioration of paint and coatings can manifest itself as peeling, chipping, cracking, ultraviolet (UV) damage and oxidation. Beyond aesthetics, these issues can affect local aerodynamics and corrosion protection. Peeling can range from a few square centimetres to more extensive areas if not addressed. The typical coating thickness is 0.1–0.5 mm [38]. Factors such as UV exposure, temperature cycling and mechanical stress can lead to crack formation. According to Ref. (39), routine visual inspections aim to detect cracks larger than 12.7 mm.
- **Lightning strikes (attachment marks).** Lightning strike attachment points are typically present as small recast/solidified metal spots, approximately 1–10 mm in diameter [40]. They are highly localised around the attachment point – often near fasteners, edges or protrusions – and are present as discrete high-contrast specks on painted aluminium or composite skins.

Consolidating these scales, an optical system capable of detecting defects in the 10–50 mm range is a valuable aid for general visual inspections. In defect-detection mode the aircraft is static, so velocity is not a factor. Such inspections are performed at scheduled intervals and can also be triggered when a specific issue is suspected [41]. Spatially, the focus is on the footprint of the aircraft (shown in green), with a light-green tolerance margin of 0.5 m around it to accommodate small placement misalignments in the bay (Fig. 3). In height, the region of interest is approximately 4.0–6.5 m, covering the wings and the upper fuselage of a narrow-body aircraft.

#### 2.4. Comparative analysis of the enabling technologies

This section brings together the assessments of MoCap, UWB and CMC systems relying on cameras discussed earlier into a streamlined comparison matrix. Table 2 contrasts their fundamental abilities, infrastructure requirements and specific trade-offs, providing a clear overview of the performance of each technology within the confines of a large aircraft hangar.

**Table 2.** Key characteristics of the three localisation and monitoring systems

Evaluation aspect	MoCap	UWB	Camera-based
Typical accuracy and rate	Sub-millimetre static ( $\approx 0.15$ mm) and $< 2$ mm dynamic errors at high update rates	Centimetre-level position when conditions are good	Varies with algorithm; not yet benchmarked, still a research topic
Core strengths	<ul style="list-style-type: none"> <li>• Highest precision available</li> <li>• Provides full 6-DoF pose</li> </ul>	<ul style="list-style-type: none"> <li>• Robust to occlusion and lighting changes</li> <li>• Low on-board processing load</li> </ul>	<ul style="list-style-type: none"> <li>• Hardware is simple (PoE cameras, switch, computing)</li> <li>• Highly versatile – same hardware can localise robots, monitor assets or inspect defects by changing software</li> </ul>
Main limitations	<ul style="list-style-type: none"> <li>• Mature commercial ecosystems</li> <li>• Very expensive hardware</li> <li>• Needs clear line-of-sight to <math>\geq 2</math> cameras per marker</li> <li>• Occlusion and wide-area calibration are challenging in hangars</li> </ul>	<ul style="list-style-type: none"> <li>• Proven in <math>&gt; 1,500</math> m<sup>2</sup> industrial halls</li> <li>• Position error rises in strong multipath and metal-rich zones; careful calibration needed</li> <li>• Extra anchors could raise complexity as the area grows</li> <li>• Not yet adopted or standardised; real-hangar robustness still unproven</li> </ul>	<ul style="list-style-type: none"> <li>• Performance is highly sensitive to illumination and viewpoint</li> <li>• Heavy computation for deep-learning pipelines</li> </ul>
Infrastructure footprint	IR cameras synchronisation, network wiring, reflective markers attached to the object of interest and calibration rigs	Fixed ceiling/wall anchors, network wiring, clock-sync links and tags on every asset	Ceiling-mounted industrial cameras, network wiring.
Scalability inside a large hangar	Precision typically drops when cameras are far apart; occlusion grows with wings and tooling	Good coverage range, but anchor count and calibration effort rise with bay size	Theoretical scalability is high, yet coverage planning and algorithm speed must align with the camera count

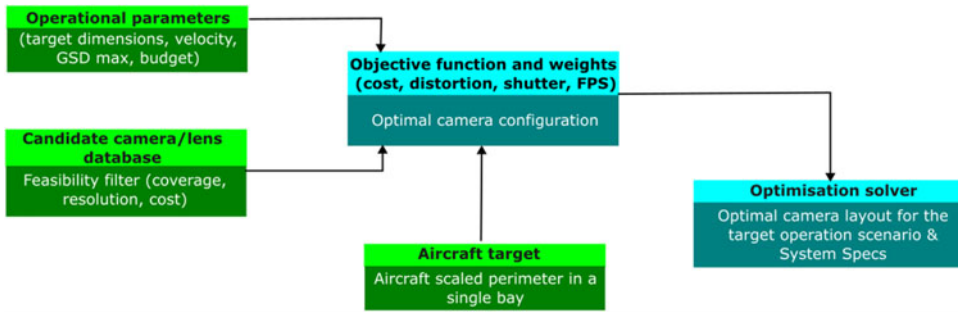


Figure 5. A high-level block diagram of the two-stage optimisation framework for camera selection and placement algorithm.

### 3. Optimised and Cost-Effective Camera Selection and Coverage Mapping

The preceding section discussed three potential systems for the smart hangar, namely MoCap, UWB and camera-based solutions. Although the first two options are available on the market, the camera-based system remains an area of open research. This section formalises the algorithmic framework for the camera-based architecture. It addresses how to select a commercially available camera-lens combination and strategically position a minimal number of units to capture the entire area of interest with the necessary ground sampling distance, frame rate and budget, depending on the operational mode (robot localisation, asset monitoring or defect detection). The ensuing framework integrates classical pinhole projection, cost-weighted distortion heuristics and binary set-cover optimisation to translate high-level operational needs into a practical camera deployment layout.

The process is divided into four steps. Initially, the envelope of the target object, the working distance and the maximum allowable ground sampling distance (GSD) are used to sift through a database of camera–lens pairs for initial feasibility. Second, an objective function balances hardware cost against optical distortion, shutter type, and frame-rate penalties, allowing each surviving pair to be ranked. In the third step, the field-of-view of the chosen pair is projected onto the hangar top-view, where a discretised grid of potential camera centres is created, and a Boolean visibility matrix is constructed. Finally, a mixed-integer linear approach addresses the resulting set-cover problem, determining both the number and positions of cameras needed to achieve full coverage with a specified overlap. Figure 5 depicts a high-level block diagram of the solution.

At first, the algorithm identifies a feasible camera-lens pair to cover a target area at a specified working distance while respecting a requested GSD, budget constraints and an approximate lens distortion penalty. This approach follows the classic pinhole camera model in photogrammetry and computer vision. The premise involves using a singular camera lens setup to observe a flat surface located at a working distance of  $d$  (in mm). The camera sensor has physical dimensions ( $S_w, S_h$ ) (in mm), and the lens has a focal length  $f$  (in mm). Under the pinhole camera approximation, the field of view (FoV) in each dimension can be calculated by projecting the sensor size onto the object plane at a distance  $d$ :

$$\text{FoV}_w = \frac{d \times S_w}{f}, \quad \text{FoV}_h = \frac{d \times S_h}{f} \tag{1}$$

The diagonal FoV is defined as:

$$\text{FoV}_{\text{diag}} = \sqrt{(\text{FoV}_w)^2 + (\text{FoV}_h)^2} \tag{2}$$

In photogrammetry, GSD represents the real-world size of one pixel in the captured image. If the camera resolution in the horizontal direction is  $R_w$  pixels (width) and  $R_h$  pixels (height), the GSD in each dimension is:

$$\text{GSD}_w = \frac{\text{FoV}_w}{R_w}, \quad \text{GSD}_h = \frac{\text{FoV}_h}{R_h} \tag{3}$$

To maintain adequate spatial resolution in the images, these values should remain below a user-defined maximum GSD limit, denoted as  $GSD_{\max}$ . Although the camera specifications lack distortion coefficients, it remains possible to implement a distortion minimisation strategy using estimations and heuristic methods. An uncomplicated method is to presume that larger fields of view result in increased distortion. The corresponding equation is as follows:

$$D = \frac{FoV_{\text{diag}}}{f} \quad (4)$$

A larger diagonal FoV for a given focal length  $f$  implies a higher distortion metric  $D$ . Although simplistic, this metric provides a convenient scalar measure to compare different camera lens configurations. The best combination of camera lenses for monitoring a target area is determined using a minimal optimisation algorithm. This algorithm considers factors such as working distance, desired coverage area, pixel density and budget constraints. The primary objective is to reduce overall costs while meeting all operational needs. Assuming that the camera cost is  $C_{\text{cam}}$  and the lens cost is  $C_{\text{lens}}$ , the definition of the total cost is the following.

$$C_{\text{total}} = C_{\text{cam}} + C_{\text{lens}} \quad (5)$$

To select among feasible camera-lens pairs, a potential multi-objective function can be:

$$O = C_{\text{total}} + \alpha D - \beta \cdot \text{GS} + \gamma \cdot \text{FPSpenalty} \quad (6)$$

where GS denotes the presence of a global shutter (binary variable), and FPS refers to the frame-per-second rate of the camera. Using this approach, the aim is to rank feasible options by combining hardware cost, distortion, shutter type preference and frame rate stability. The constants  $\alpha$ ,  $\beta$  and  $\gamma$  allow tuning of trade-offs between distortion, shutter preference and temporal resolution. More specifically, the parameter  $\alpha$  is a user-chosen weight to balance monetary cost versus the distortion penalty  $D$ . In practice, the parameter  $\alpha$  may be chosen according to the desired level of penalisation for wide fields of view or lens distortion. The parameter  $\beta$  assigns a cost-reduction bonus to configurations that include a global shutter sensor. This reflects the practical advantage of global shutters in avoiding motion blur during dynamic operations, making them more suitable for applications involving moving objects. The parameter  $\gamma$  penalises configurations with non-ideal frame rates. Frame rates below 20 FPS may result in choppy visual streams, while rates above 50 FPS are often unnecessary and can increase bandwidth and processing load. This term helps to prioritise configurations that operate within a desirable temporal resolution range.

Before applying ranking configurations, a filtering step is applied to discard those that do not meet these fundamental criteria. This ensures that only operationally feasible combinations are passed on to the optimisation process, reducing computational overhead and improving the quality of the result.

In terms of constraints, each candidate configuration must satisfy the following:

- Coverage:

$$FoV_w \geq W_{\text{target}}, \quad FoV_h \geq H_{\text{target}} \quad (7)$$

where  $W_{\text{target}}$  and  $H_{\text{target}}$  is the required coverage size (in mm).

- Resolution:

$$GSD_w \leq GSD_{\max}, \quad GSD_h \leq GSD_{\max} \quad (8)$$

- Budget:

$$C_{\text{total}} \leq C_{\max} \quad (9)$$

where  $C_{\max}$  is the maximum allowable budget (in British pounds).

Once the desired dimensions for the camera sensor have been determined, the next step is to estimate the FoV of the camera positioned on the ceiling to inspect the aircraft. A geometric model was utilised,

taking into account the sensor dimensions and the lens’s focal length. The angular field of view (in degrees) is calculated for both the width and height of the camera sensor using the following equation:

$$\text{FoV} = 2 \cdot \tan^{-1} \left( \frac{s}{2f} \right) \cdot \left( \frac{180}{\pi} \right) \tag{10}$$

where

- $s$  is the sensor dimension (width or height) in millimetres
- $f$  is the focal length in millimetres

During the space allocation phase, accurately scaling the aircraft’s shape to real dimensions is crucial. The two most popular original equipment manufacturers (OEMs) provide online catalogues of their aircraft models, commonly called 3-view drawings, which are used for gate planning purposes [42, 43]. The drawings are provided in the Drawing eXchange Format (DXF) and include the top, side, and front perspectives. However, the presentation of these drawings exhibits inconsistency, both between different OEMs but also between different models produced by the same manufacturer.

The suggested algorithm necessitates that the aircraft’s top view is represented as a closed polygonal line. To form this polygon, the procedure involves importing the 3-view drawing into a vector drawing application, that supports Scalable Vector Graphics (SVG) and manually outlining the perimeter polygon on a new layer above the image. Since this task must be performed only once per aircraft model, it is not considered a heavy or unsustainable task within the suggested process flow.

Once the exterior representation of a narrow-body aircraft is obtained in SVG format, it is parsed, and an ordered vertex list is created.

$$\mathcal{V} = \{(x_i, y_i)\}_{i=1}^N \tag{11}$$

with  $N$  the number of vertices. Assuming that  $L_{\text{px}}$  is the pixel length of the fuselage,

$$L_{\text{px}} = \max_i x_i - \min_i x_i, \tag{12}$$

and let  $L_m$  be the length of the aircraft. The metric scale factor is therefore

$$\alpha = \frac{L_m}{L_{\text{px}}}, \quad (x_i^*, y_i^*) = (\alpha x_i, \alpha y_i). \tag{13}$$

The reference polygon (Fig. 9(a)) can be defined by the scaled vertex as

$$\mathcal{P} = \text{poly} \{(x_i^*, y_i^*)\} \subset \mathbb{R}^2. \tag{13}$$

To guarantee safe stand-off for cameras or for the aircraft model, an external inspection envelope was introduced,  $\mathcal{P}$  is uniformly offset by a Minkowski sum with radius  $\delta$ :

$$\mathcal{P}^+ = \{\mathbf{p} \in \mathbb{R}^2 | \text{dist}(\mathbf{p}, \mathcal{P}) \leq \delta\}. \tag{14}$$

This is like inflating the perimeter of the plane by a specific distance (Fig. 9(b)). In the next stage, a discretisation of the region of interest occurs. A regular square grid of spacing

$$\Delta = \lambda \tag{15}$$

is seeded over the bounding rectangle of  $\mathcal{P}^+$ . The parameter  $\Delta$  defines the granularity, but also affects the computation time of the approach. Let

$$\mathbb{G}_\Delta = \{(x, y) \in \mathbb{R}^2 | x = x_{\min} + k\Delta, \quad y = y_{\min} + \ell\Delta\} \tag{16}$$

be that lattice. The set of target points is determined by a Boolean flag based on whether interest is directed toward the internal (Fig. 9(c)) or external space (Fig. 9(d)) of the aircraft, depending on the specific scenario.

$$\mathbf{G} = \begin{cases} \{g \in \mathbb{G}_\Delta | g \in \mathcal{P}^+\}, & \text{internal coverage} \\ \{g \in \mathbb{G}_\Delta | g \notin \mathcal{P}^+\}, & \text{external coverage} \end{cases} \tag{17}$$

The cameras are mounted at a hangar ceiling height of  $h$  and possess horizontal and vertical fields of view. In a pinhole model, the ground footprint of one camera is the rectangle  $[-\frac{W}{2}, \frac{W}{2}] \times [-\frac{L}{2}, \frac{L}{2}]$  with

$$W = 2h \tan(\theta_h/2), \quad L = 2h \tan(\theta_v/2). \tag{18}$$

Let each ceiling-mounted camera project a ground-plane footprint of width  $W$  and length  $L$

$$\Delta_x = \beta W, \quad \Delta_y = \beta L. \tag{19}$$

The parameter  $\beta \in (0, 1]$  is defined as overlap ratio. This corresponds to a ground-plane footprint overlap of  $(1 - \beta)$  in each axis, or equivalently,  $(1 - \beta) \times 100\%$  overlap. For example,  $\beta = 0.80$  produces an overlap of 20% between adjacent camera footprints. In this work, the overlap defines the region of shared ground coverage used to support cross-camera handover and track continuity in localisation scenarios. The same formulation allows reduced overlap in static inspection tasks, where object handover is unnecessary.

The origin of this lattice is shifted so that it coincides with the centroid of  $\mathcal{P}^+$ , ensuring an approximately symmetric layout around the aircraft. Let  $\{\mathbf{c}_1, \dots, \mathbf{c}_m\}$  denote all candidate locations generated within the enlarged boundary box. Then, for every point on the grid  $\mathbf{p}_i \in \mathbf{G}$  and camera candidate  $\mathbf{c}_j$ , an indicator was defined as

$$A_{ij} = \begin{cases} 1, & \mathbf{p}_i \in \text{Cov}(\mathbf{c}_j), \\ 0, & \text{otherwise,} \end{cases} \tag{19}$$

where  $\text{Cov}(\mathbf{c}_j) = \mathbf{c}_j + [-\frac{W}{2}, \frac{W}{2}] \times [-\frac{L}{2}, \frac{L}{2}]$ . At this stage, the Boolean matrix  $\mathbf{A} \in \{0, 1\}^{n \times m}$  compactly describes which cameras see which points. A binary decision variable  $x_j \in \{0, 1\}$  indicate whether the camera  $j$  is installed. Ultimately, the minimum-camera problem is the classical set cover formulated as a binary integer linear programme:


$$\min_{\mathbf{x} \in \{0,1\}^m} \sum_{j=1}^m x_j \tag{20}$$

$$\text{subject to } \sum_{j=1}^m A_{ij} x_j \geq 1, \quad i = 1, \dots, n. \tag{21}$$

Equations (20)–(21) are implemented using the PuLP [44] library with category LpBinary (Linear-programming of binary type) and solved using the CBC MILP (COIN-OR branch-and-cut – mixed-integer linear programming) solver engine [45]. The optimal index set  $\mathcal{C}^* = \{j | x_j^* = 1\}$  specifies both the number and the positions of cameras (Figs. 10–15) required for full coverage of the target region defined in Equation (17).

#### 4. Design-to-Cost Deployment Blueprints

Automated robotic inspection forms an integral part of the smart hangar concept, which utilises Industry 4.0 technologies to augment the abilities of human experts to perform necessary maintenance of the aircraft. The primary focus is on commercial aircraft that can be divided into two major categories. Narrow-body aircraft, also known as single-aisle aircraft, are the workhorses of short- to medium-haul routes. They typically have a single aisle, two or three seats on each side of the aisle, a capacity of 100–240 passengers, and a range of 3,000–4,000 nautical miles. Examples include Boeing 737 series and Airbus A320 family (A318, A319, A320, A321). The second category is the wide-body aircraft, or twin-aisle aircraft, which are designed for long-haul flights and have two aisles, seven to ten seats in economy class, capacity for 200–550 passengers, and a range of 5,000–8,000 nautical miles. Examples include Boeing 767, 777, 787 Dreamliner and Airbus A330, A350 and A380. Figure 6 illustrates the difference in size between the narrow-body and wide-body aircraft.



**Figure 6.** The size comparison between narrow-body (Airbus A320) and wide-body (A380) aircraft.


Taking into account the analysis presented in Ref. (46) for the distribution of the fleet, the global operating commercial aircraft fleet in 2025 is dominated by narrow-body jets, which account for approximately 62% of the total. Wide-body, regional and turboprop jets are projected to comprise the remaining 38% of the worldwide commercial aircraft fleet.

The hangar dimensions are adjusted to the aircraft dimensions. For example, a hangar intended to host only narrow-body aircraft is required to have a door height of 14 m and an opening of 37 m. In the case of wide-body, the specifications are 24 m in height and 75 m opening. In terms of ceiling height, the first case requires reaching 18 to 20 m, while the second case requires reaching 26–28 m. In addition, the hangars can accommodate more than one aircraft simultaneously, featuring dedicated, predefined spaces called bays. The typical lateral dimensions [35] for a narrow-body aircraft hangar range from  $40 \times 50$  m to  $60 \times 60$  m.

In the following sections, potential architectures are provided for each case. In the MoCap and UWB options, manufacturers provide suggestions for system deployments with an order-of-magnitude cost associated. For the camera-based system, a technical analysis is presented, based on the methodology outlined earlier, along with an estimated cost based on the components.

#### 4.1. MoCap blueprint

The MoCap blueprint is a configuration proposed by a lead manufacturer, tailored specifically for a single-bay narrow-body hangar. The proposal was parameterised by the dimensions of bay width, length and ceiling height. The camera configuration involves installing 12 cameras at a height of 22 m, pointing downward (Fig. 7(a)). Side or low-position cameras are excluded, as they are impractical for hangars with multiple bays. Local occlusions at the wing–fuselage junctions are the dominant risk. Consequently, additional cameras can improve coverage in these regions. This contingency also accounts for operational occlusions (e.g., lifts, tugs, tool trolleys) that are absent from the vendor’s aircraft-only simulation and are expected to be encountered in a real service environment. The 12 cameras effectively cover a distance of about 32 m for a 19-mm marker. The suggested model is a 12MP (300 fps) camera, with a maximum capturing distance of 40 m and a 3D resolution of 0.04 mm. Figure 7(b) shows the capture volume, which demands that each marker be captured by three cameras. In reality, the system can track markers over a broader area than this, although a lower marker resolution might result in occasional tracking loss. In a MoCap system, tracking each marker requires the use of at least two cameras.



**Figure 7.** Top (a): MoCap system with 12 cameras set-up. Bottom left (b): MoCap system with 3-camera minimum tracking for each marker. Capture volume in shaded green. Bottom right (c): MoCap system with 2-camera minimum tracking for each marker.

Figure 7(c) illustrates the capture volume with two-camera tracking, where a larger area is covered toward the nose and tail. In practice, increasing the range of the camera does not provide much benefit. In addition, the size of the marker and the camera range increase proportionally, and markers of 30, 40 and 50 mm are commercially available, depending on the size of the robotic platform selected for the inspection. For hangar-scale volumes, vendors increasingly replace dynamic wanding with fixed-camera calibration referenced to a surveyed control network. In practice, a small set of permanently mounted markers is measured with a total station and entered as control points to the MoCap tracking manager, allowing the solver to recover extrinsics and absolute scale against these surveyed points [47]. Independent validations that tie MoCap volumes to geodetic references further support this workflow for large industrial spaces. The projected expense for this system was assessed to fall within the range of £180,000 to £200,000, as stated in the quotation provided by a leading motion capture manufacturer in April 2025.

#### 4.2. UWB blueprint

This blueprint reflects a manufacturer-validated UWB configuration tailored to a single-bay narrowbody hangar, parameterised by bay width, length and ceiling height. Within the UWB coverage of the hangar, the extent of the necessary infrastructure is determined by the specific tracking applications used and

the efficiency required in regions with significant obstructions, such as under the wings. In scenarios involving localisation and monitoring, such as the tracking of drones and ground platforms, the process is considerably simpler compared to tracking certain assets (e.g., tools). This is because it typically occurs above or between highly obstructed areas, avoiding locations like under wings. Considering the initial assumption of a single bay maintenance zone sized between  $40 \times 50$  m and  $60 \times 60$  m (resulting in a volume from 2,000 to 3,600 m<sup>3</sup>) and noting that the application occurs in fairly open settings, a prudent estimate is to account for 10–25 anchors to achieve adequate coverage. For comparison, when applications involve tracking under the wing and fuselage, and around the engines, more sensors, possibly ranging from 20 to 40, might be used.


On the implementation side, power and network connectivity are required for the installation of each anchor. The most common practice involves using a Cat5e UTP Ethernet cable to connect each sensor location from a control room, which needs to be within 100 m of the sensor mounting location to comply with Ethernet networking limits. Given the size of the hangar, this likely necessitates equipment placement in two control rooms, each equipped with a timing distribution unit (TDU) for precise time synchronisation of the anchors and a PoE networking switch to provide power and networking to each sensor. The cost associated with running each cable can vary significantly depending on the location. It is largely determined by whether the cabling is installed during the construction of the hangar or after, whether the installation after construction occurs while the space is active (it is considerably more expensive if the space is active) and the geographical location where this activity occurs. Lastly, the cost of actually mounting the sensor on its bracket at the end of the cable needs to be included.

In order for the system to be completely functional, every object of importance must have a tag installed. A wide range of options are available regarding size, power consumption and update rate. As an illustration, a small tag may feature dimensions of  $46 \times 42 \times 18$  mm and weigh about 20 g, with a battery life that exceeds ten years while transmitting continuously at an update rate of 1 Hz. For drone tracking applications, the tags can be configured to transmit at a higher rate (e.g. 10 Hz), wherein the battery life might be approximately 12–15 months. These tags are priced at about £40 each. Another option is to incorporate a tag module into the drone or ground platform. This module draws its power from an external source, such as the drone's power system, eliminating the need for an internal battery that could run out. Naturally, the drone itself is periodically recharged. Without a battery restriction, the tag module can achieve considerably higher update rates, such as 50 Hz.

The UWB infrastructure is capable of monitoring thousands of tags within its coverage area and can deliver aggregate update rates of several thousand updates each second. The resulting data stream, composed of (ID,  $x$ ,  $y$ ,  $z$ , timestamp), is available for transmission to other applications. Although system suppliers offer various applications, there is also the possibility of transmitting the data as UDP network data, allowing it to be used by third-party connected systems. According to a quotation from a leading UWB manufacturer issued in May 2025, the configuration of a candidate UWB system to cover a single bay hangar, which encompasses anchors, PoE switches, TDUs, wiring and calibration, is projected to cost approximately £49,000. This estimate is a general figure that may vary depending on the number of anchors used. An example setup of a UWB system is installed at the Digital Aviation Research and Technology Centre (DARTeC) at Cranfield University, where nine anchors and two types of tags with 0.5 and 30 Hz frequencies have been installed (Fig. 8).

#### 4.3. Ceiling-mounted camera system blueprint

The camera-based system is not a commercially available product, so the candidate implementation is not based on the suggestion of any manufacturer. Depending on the application scenario, the algorithms presented in Section 3 suggest a candidate system that can perform the robot localisation, monitoring or defect detection task. In the beginning, extensive product research was performed in order to identify the commercially available products that could be used as data for the analytical approaches. The analysis is focused on industrial machine vision cameras, excluding surveillance or CCTV devices. The automatic adjustments and compressed video streams of surveillance cameras impede precise control over GSD, exposure, timing, and synchronisation, which are essential to achieve reproducible and pixel-level performance.



**Figure 8.** *The UWB system installed at Cranfield's DARTeC smart hangar.*

In system design, camera specifications are crucial for optimal system performance. Beyond selecting the right camera and lens combination, it is essential to consider cost-effectiveness and the need to depict the object of interest at a suitable size for the deep learning component. The chosen camera features that were used as a criterion for the analysis include sensor size and format, resolution, shutter type, pixel dimensions, frame rate, connectivity speed and cost (Table 3). For the lens, the most important characteristic was the focal length (Table 4). The compatibility between the camera and the lens was ensured as both supports the C-mount interface.

The machine vision cameras support two connection interfaces: USB 3.0 and PoE. The PoE interface is standardised because it enables both power and data to be transmitted through a single Cat6 cable for distances up to 100 m, facilitating centralised switching and scalable aggregation on a hangar scale. In contrast, ceiling-side compute is forced by USB 3.0 short-reach cabling (or active extenders), making it less suitable for robust hangar installations.

The typical velocity of the target object is associated with the camera frame rate to generate a suitable sequence of frames to track movement. The linked factor that influences system performance is the GigE connectivity interface, which is related to the data transfer rate, latency and overall system throughput. The maximum theoretical data transfer rate of a standard GigE interface is 1 gigabit per second, which is equal to 125 megabytes per second. This determines the maximum frame rate and resolution that a camera can output. So, for example, a 2MP camera outputs uncompressed 8-bit images at 60 fps  $\times$  120 MB/s, which is close to the GigE limit. In general, GigE has higher latency compared to USB 3.0 interfaces.

The final attribute evaluated was the shutter type. The global shutter camera takes a complete image in a single instance, as every pixel is exposed at the same time. There is no motion distortion, so fast-moving objects are captured without skew or wobble. This is the ideal case for precise measurements, which is useful in robotics, industrial vision and 3D reconstruction. They can support easier synchronisation with external sensors like LiDAR or IMU, because the timestamp is consistent across the whole image. However, the rolling-shutter camera captures the image line-by-line from top to bottom (or side to side), over a short time interval. Because of this mode of operation, the image contains motion artefacts, which can cause skew, wobble or partial exposure when objects move quickly or the camera moves during capture. In addition, it is harder to sync since different parts of the image correspond to slightly different times.

The video streams from the cameras are transmitted to the computing unit through an Ethernet switch. To facilitate a simpler installation and reduce potential failure points, a PoE switch is suggested instead of using the power supply for each camera separately. When using high-resolution GigE cameras, it is

**Table 3.** The table presents a large selection of PoE cameras. The information was gathered from the Edmund Optics website, and the prices were captured on the 30th of March 2025

Brand	Sensor (mm)	Resolution (px)	Format (inches)	Mpixels	Shutter	Pixel size ( $\mu\text{m}$ )	Frame rate	GigE	Cost (£)
Basler(1)	11.25×7.03	1,920×1,200	1/1.2"	2.3	Global	5.86	42.0	1.0	535
Basler(2)	9.22×5.76	1,920×1,200	2/3"	2.3	Global	4.8	50.0	1.0	612
Basler(3)	11.25×7.03	1,920×1,200	1/1.2"	2.3	Global	5.68	50.0	1.0	1,032
Basler(4)	7.07×5.3	2,048×1,536	1/1.8"	3	Global	3.45	36.0	1.0	522
Basler(5)	11.26×11.26	2,048×2,028	1"	4.2	Global	5.5	25.0	1.0	1,989
Basler(6)	5.7×4.28	2,592×1,944	1/2.5"	5	Rolling	2.2	14.0	1.0	497
Basler(7)	12.43×9.83	2,590×2,048	1"	5	Global	4.8	5.0	1.0	943
Basler(8)	6.44×4.62	3,856×2,764	1/2.3"	10.6	Rolling	1.67	10.0	1.0	544
Basler(9)	7.41×4.95	3,088×2,064	1/1.8"	6	Rolling	2.4	16.0	1.0	385
Basler(10)	7.44×5.62	4,024×3,036	1/1.7"	12	Rolling	1.85	8.0	1.0	463
Basler(11)	14.13×7.45	4,096×2,160	1"	9	Global	3.45	12.0	1.0	1,487
Basler(12)	14.13×10.35	4,096×3,000	1.1"	12	Global	3.45	8.0	1.0	1,819
Basler(13)	13.13×8.76	5,472×3,648	1"	20	Rolling	2.4	5.0	1.0	654
Basler(14)	6.6×4.1	1,920×1,200	1/2.3"	2.3	Global	3.45	51.0	1.0	332
Basler(14)	5.2×3.9	2,592×1,944	1/2.8"	5	Rolling	2.0	22.0	1.0	279
Basler(15)	7.7×4.3	3,840×2,160	1/1.8"	8.3	Rolling	2.0	13.0	1.0	414
Basler(16)	11.2×8.2	4,096×3,000	1/1.1"	12.3	Global	2.74	44.0	1.0	1,827
Basler(17)	14.6×8.3	5,320×3,032	1.1"	16.1	Global	2.74	34.0	1.0	1,989
Basler(18)	12.3×12.3	4,504×4,504	1.1"	20.2	Global	2.74	27.0	1.0	2,329
Basler(19)	14.6×12.6	5,328×4,608	1.2"	24.4	Global	2.74	22.0	1.0	2,851
Lucid(1)	7.37×4.92	3,072×2,048	1/1.8"	6.3	Rolling	2.4	42.8	2.5	436
Lucid(2)	14.44×9.9	3,208×2,200	1.1"	7.1	Global	4.5	35.0	2.5	1,270
Lucid(3)	8.57×7.17	2,448×2,048	2/3"	5.0	Global	3.45	35.6	2.5	697
Lucid(4)	7.78×7.78	2,840×2,840	2/3"	8.1	Global	2.74	31.8	2.5	901
Lucid(5)	14.13×7.45	4,096×2,160	1"	8.9	Global	3.45	13.7	1.0	1,041
Lucid(6)	14.58×8.31	5,320×3,032	1.1"	16.2	Global	2.74	17.5	2.5	1,445

*Table 3. Continued*

Brand	Sensor (mm)	Resolution (px)	Format (inches)	Mpixels	Shutter	Pixel size ( $\mu\text{m}$ )	Frame rate	GigE	Cost (£)
Lucid(7)	14.13×10.35	4,096×3,000	1.1''	12.3	Global	3.45	22.6	2.5	1,627
Lucid(8)	11.22×8.22	4,096×3,000	1/1.1''	13.2	Global	2.74	21.7	2.5	1,096
Lucid(9)	13.13×8.76	5,472×3,648	1''	20	Rolling	2.4	6.0	1.0	522
Lucid(10)	13.13×8.76	5,472×3,648	1''	20	Rolling	2.4	9.0	2.5	663
Lucid(11)	12.34×12.34	4,504×4,504	1.1''	20.4	Global	2.74	13.5	2.5	1,840
Lucid(12)	14.58×12.6	5,320×4,600	1.2''	24.5	Global	2.74	11.3	2.5	2,057
Flir(1)	14.13×7.45	4,096×2,160	1''	8.9	Global	3.45	13.8	1.0	1,389
Flir(2)	14.13×10.35	4,096×3,000	1.1''	12.3	Global	3.45	9.9	1.0	1,717
Flir(3)	13.13×8.76	5,472×3,648	1''	20.0	Rolling	2.4	6.1	1.0	595
Flir(4)	14.6×12.6	5,320×460	4/3''	24.5	Global	2.74	5.0	1.0	1,891
Allied(1)	8.5×7.1	2,464×2,056	2/3''	5.1	Global	3.45	95.0	5.0	1,377
Allied(2)	7.8×7.8	2,848×2,848	2/3''	8.1	Global	2.74	59.0	5.0	1,224
Allied(3)	7.4×5.6	4,024×3,036	1/1.1''	12.4	Global	2.74	39.0	5.0	1,627
Allied(4)	7.44×6.61	4,024×3,036	1/1.7''	12.2	Global	1.85	41.0	5.0	760
Allied(5)	14.59×8.33	5,328×3,040	1.1''	16.2	Global	2.74	30.0	5.0	1,899
Allied(6)	12.36×12.36	4,512×4,512	1.1''	20.4	Global	2.74	24.0	5.0	2,231
Allied(7)	13.19×8.81	5,496×3,672	1''	20.2	Rolling	2.4	25.0	5.0	1,024
Allied(8)	14.6×12.63	5,328×4,608	1.2''	24.6	Global	2.74	20.0	5.0	2,677

**Table 4.** The table presents a selection of lenses (C-mount) compatible with the cameras presented in Table 3. The information was gathered from the Edmund Optics website, and the prices were captured on the 30th of March 2025

Brand	Type	Focal length (mm)	Cost (£)
Techspec(0)	4mm UC Series Fixed Focal Length Lens	4	208
Techspec(1)	6mm UC Series Fixed Focal Length Lens	6	221
Techspec(2)	8mm UC Series Fixed Focal Length Lens	8	221
Techspec(3)	12mm UC Series Fixed Focal Length Lens	12	221
Techspec(4)	18mm UC Series Fixed Focal Length Lens	18	518
Techspec(5)	25mm UC Series Fixed Focal Length Lens	25	518
Techspec(6)	35mm UC Series Fixed Focal Length Lens	35	518
Techspec(7)	50mm UC Series Fixed Focal Length Lens	50	518
Techspec(8)	75mm UC Series Fixed Focal Length Lens	75	561
Techspec(9)	100mm UC Series Fixed Focal Length Lens	100	561

**Table 5.** The table presents the camera-lens optimal combination per target scenario

Specifications	Scenario A <sup>1</sup>	Scenario B <sup>2</sup>	Scenario B <sup>2</sup>	Scenario C <sup>3</sup>	Scenario C <sup>3</sup>
Target dimensions (mm)	40 × 40	500 × 500	800 × 500	2,400 × 1,100	600 × 450
Target dimensions (px)	45 × 45	83 × 83	114 × 71	490 × 224	121 × 91
Target area (m)	3 × 3	7 × 7	14 × 14	20 × 20	15 × 15
GSD (mm/pixel)	0.89	6.02	6.94	4.91	4.94
Camera	Allied(7)	Basler(1)	Lucid(2)	Lucid(11)	Allied(3)
Lens	Techspec(7)	Techspec(4)	Techspec(3)	Techspec(3)	Techspec(2)

Notes: <sup>1</sup>Scenario A is related to mid-size defect detection on the surface of the aircraft.

<sup>2</sup>Scenario B is related to localisation of drone (column 2) and ground robotic platform (column 3).

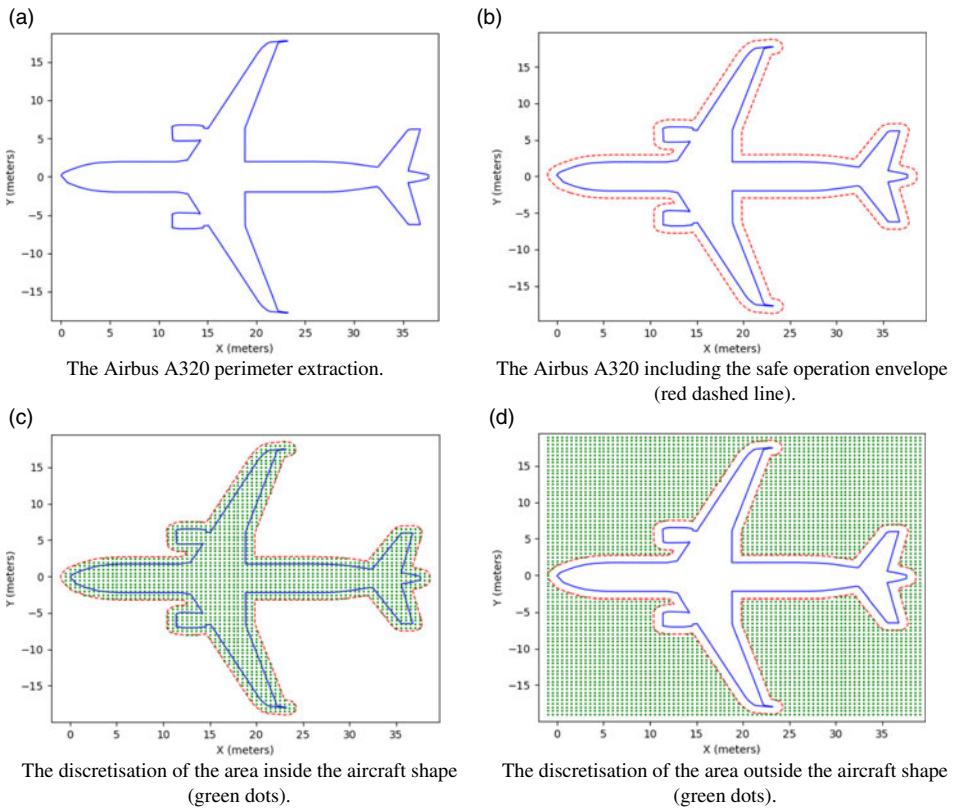
<sup>3</sup>Scenario C is related to monitoring of ground assets (column 4) and humans (column 5).

recommended that Ethernet switches and Network Interface Cards (NICs) be used to support jumbo frames. In general, in Ethernet networking, a standard frame has a Maximum Transmission Unit (MTU) of 1,500 bytes. Jumbo frames extend this limit, typically allowing up to 9,000 bytes per frame. This means that more data can be transmitted in a single packet, reducing the number of packets needed for large data transfers. The advantages of utilising Jumbo frames include:

- Reduced CPU load: Fewer packets mean fewer interrupts for the CPU to handle, lowering processing overhead.
- Improved throughput: Larger frames can carry more data, enhancing overall network efficiency.
- Decreased packet loss: With fewer packets on the network, the chance of collisions and packet loss diminishes.

There are many commercial options available that satisfy these requirements. For example, the Mikrotik cloud router switch (CRS328-24P-4S+RM), which offers 24 PoE gigabit ports, with 4 SFP (small form-factor pluggable) ports to connect fibre optics of different types and speeds. This specific device supports 10,218-byte jumbo frames. A representative cost of this device is £417 (price captured from Amazon.co.uk on the 27th of April 2025).

The candidate combinations of camera and lens are fed into the camera selection algorithm. Depending on the target scenario, the algorithm proposes five potential configurations that are illustrated in Table 5.



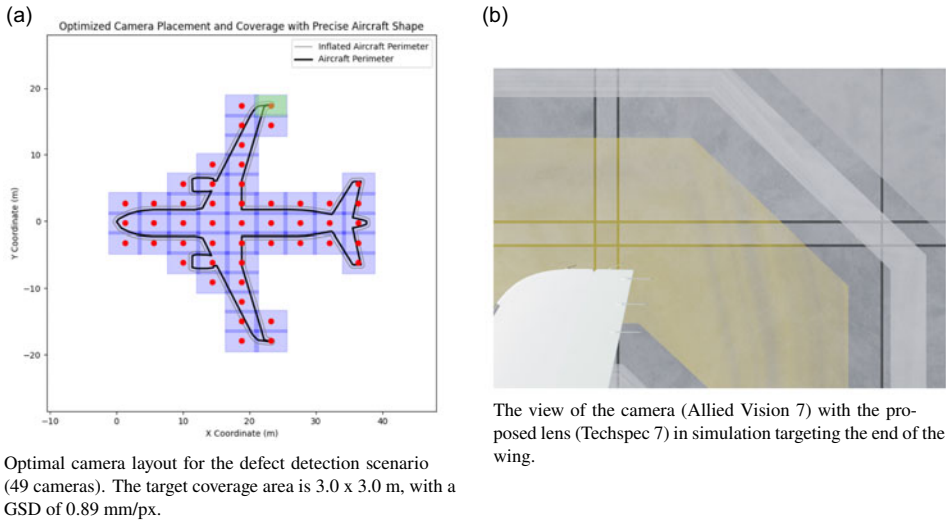
**Figure 9.** *Different candidate modes of operation for the CMC system.*

Starting with the Airbus A320 aircraft as the base case, the initial step involves obtaining the top-view perimeter from the SVG file. This outline is then scaled to match a total length of 37.6 m, in accordance with the manufacturer's publicly available specifications. The result is depicted in Fig. 9(a).

The next step involves the extension of the perimeter to create a safe operating envelope to accommodate tolerances related to the target scenario. For instance, in drone localisation, the value is 1 m to include a buffer space for the manoeuvres, but in defect detection, it is 0.5 m as there is no moving target and only aircraft misalignment should be considered. The result is illustrated in Fig. 9(b).

Once the operating envelope is determined based on the specified target scenario, the next step involves discretising the region of interest. In cases where the focus is on the aircraft, such as in drone localisation, the area is transformed into a mosaic pattern based on the set discretisation distance, as demonstrated in Figs 9(c) and (d). Experimental results demonstrated that a value of 0.5 m in grid spacing achieves a balance between solution accuracy and the necessary computational time. Because the coverage discretisation uses a square grid, the number of coverage points grows roughly with the inverse square of the spacing, so the run time increases nearly quadratically as the grid is refined.

In the final stage, the algorithm, considering the discretised region of interest of the previous step and the suggested camera field of view projection, identifies the optimal camera layout for the bay. In the process, an overlap factor was also taken into account. This factor can affect system performance under varying conditions, especially with regard to transferring object localisation between cameras. If the scenario, such as defect localisation, does not necessitate it, the overlap is reduced. This parameter may also prove vital when stitching camera feeds into a larger image, as the presence of shared salient points becomes essential.



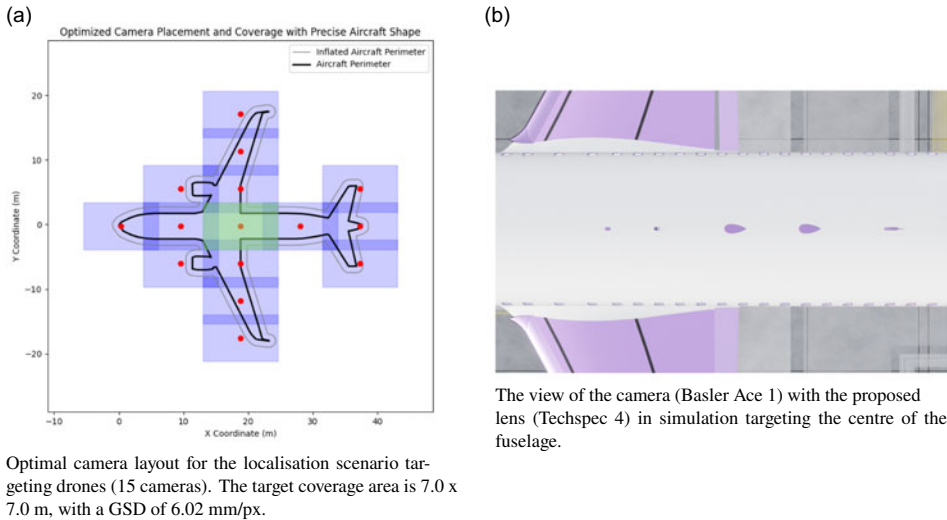
**Figure 10.** Defect detection scenario for the camera-based system.

#### 4.3.1. Scenario A – medium-defect mapping

The initial case of interest involves the detection of defects. This scenario focusses on identifying flaws present on the top surfaces of the aircraft, such as the fuselage and wings. The existing commercial systems, such as UWB and Mocap, are not designed for this scenario as they focus on localisation and tracking through the use of tags or reflective markers. This is notably intriguing because for maintenance technicians, it is safer to avoid working at heights with scaffolding and cherry pickers. Throughout the experimental process, it became apparent that effective identification of small defects using cameras installed on ceilings presents significant physical challenges. This threshold emerged from GSD-based calculations, which indicated that capturing adequate pixel density for small surface features requires extremely high-resolution sensors or a highly concentrated network of cameras. Both options greatly increase costs and complexity. In contrast, focussing on medium-sized defects of size  $40 \times 40$  mm presented a more reasonable and feasible compromise, enabling the system to remain scalable and cost-efficient while still offering significant assistance for general visual inspections. In the medium-size defect case, a potential target will occupy approximately 45 px in the captured image. Given that the targets are located roughly between 16 and 19 m away, corresponding to the aircraft's upper surface, the proposed configuration encompasses an area of 3 by 3 m. To cover the entire surface of a narrow-body aircraft, 49 cameras are required (Fig. 10). Considering the additional cost of three 24-port PoE Ethernet switches, the cost adds up to a total of £76,809. In this cost estimate, cables are excluded since the installation cable length estimate was beyond the scope of the work. However, a representative price for 100 m of unshielded twisted pair (UTP) Category 6 cable is approximately £60 (price captured from Amazon.co.uk on the 27th April 2025). For this scenario, the cameras were configured to have a 10% overlap in coverage to minimise the number of cameras needed. Moreover, as the target remains stationary, it is unnecessary to maintain a significant margin for the target transition between the cameras. Similarly, since the target is not moving, there is no negative impact that the suggested camera has a rolling shutter or the relatively low frame rate of 25 fps.

#### 4.3.2. Scenario B – localisation for robotic platforms

The significance of the localisation case is considerable, as it functions as an external system supplying the absolute position of a robotic platform, whether aerial or ground. These data can then be integrated with onboard odometry to enhance navigation reliability. Based on the dimensions of the target and the distance from the camera (ceiling), there are two potential cases: drones and ground robotic platforms.



**Figure 11.** Localisation scenario targeting drones for the camera-based system.

Depending on the robot’s operation space, the distance is approximately 15–18 m for the drones and 22 m for the ground platforms.

In this context, it is crucial to confirm that the anticipated and permissible velocity of the platforms is compatible with the frame rate of the cameras. The equations presented in 22 can be used to determine the distance travelled between frames. The camera for drones has 42 fps, and for the UGVs, 35 fps. With a maximum speed of 1.5 m/s, the distance between each frame for a 40 fps camera is 3.75 cm. In addition, both cameras suggested by the algorithm are equipped with a global shutter that is recommended for improved performance on moving targets.

$$\text{Time per frame} = \frac{1}{f}, \quad \text{Distance per frame} = v \times \text{Time per frame} \quad (22)$$

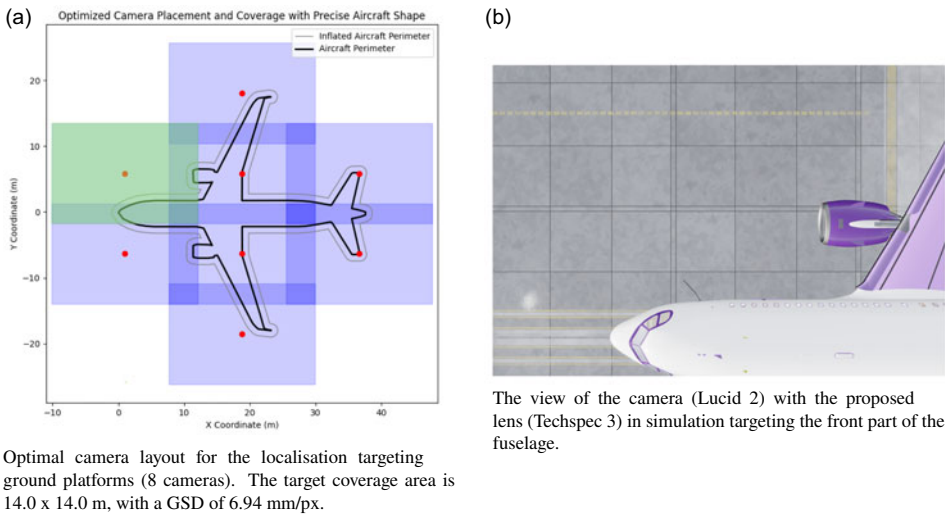
In the case of aerial platforms, Fig. 11 illustrates that covering the entire surface of a narrow-body aircraft requires 15 cameras with the specific lenses. The total cost amounts to approximately £16,500 including a 24-port PoE Ethernet switch.

The ground-plane footprint overlap is set at 20% to support cross-camera handover and track continuity for moving robots, that is, to maintain association as targets move between adjacent FoVs. In order to contextualise the chosen overlap ratio within the localisation scenario, the duration of dual visibility available for cross-camera handover is measured by the observed ground-plane fields of view and standard operating conditions.

Assuming that the horizontal (side-to-side) and vertical (longitudinal) footprints of each ceiling-mounted camera are  $W = 11.562$  m and  $L = 7.226$  m, respectively, as calculated for the selected camera and lens. With an overlap ratio  $r = 0.20$ , the overlap zones are  $d_{ov,x} = rW = 2.3124$  m (horizontal adjacency) and  $d_{ov,y} = rL = 1.4452$  m (vertical adjacency). For a drone with a typical lateral size  $s = 0.5$  m moving at  $v = 1$  m/s, the conservative full-target co-visibility lane is  $d_{eff} = \max(d_{ov} - s, 0)$ . Given a frame rate  $f = 42$  fps, based on the selected camera model, the available association time and frame count during handover are

$$t = \frac{d_{eff}}{v}, \quad N = ft \quad (23)$$

Numerically, this yields  $d_{eff,x} = 1.8124$  m,  $t_x = 1.8124$  s, and  $N_x \approx 76$  frames for horizontal handover; and  $d_{eff,y} = 0.9452$  m,  $t_y = 0.9452$  s, and  $N_y \approx 40$  frames for vertical handover. Therefore, even in the more challenging vertical case (requiring the entire target 0.5 m to remain within both footprints), the



**Figure 12.** Localisation scenario targeting ground robotic platforms for the camera-based system.


system provides approximately 40 frames of dual visibility in 1 m/s, which is sufficient for robust cross-camera association and handover. This scenario-grounded calculation motivates the use of  $r = 0.20$  as it provides reliable handover while avoiding the increase of cameras that higher overlaps entail.

In the context of the localisation of the ground robotic platform, the demands for camera specifications pertinent to object detection can be eased due to the increased dimensions, resulting in the installation of eight cameras (Fig. 12). The estimated cost, including a 24-port PoE Ethernet switch, is approximately £12,800. For ground robots, the overlap setting is retained. Due to their wider lateral reach, slower rotational speeds and the planar nature of their trajectories constrained by the ground floor, the pixel size requirement per target is less stringent compared to drones. From an overhead camera viewpoint, the drone's X-shaped open frame and central payload create a visually fragmented target that poses challenges for segmentation, particularly against the bright, reflective surfaces of the aircraft it inspects. In contrast, ground platforms have solid bodies and traverse the darker hangar floor, providing higher contrast and more stable silhouettes.

Using drones for inspections near the target surface offers an effective solution to the challenges associated with defect detection, which requires multiple cameras, to identify medium to large defects as outlined in Scenario A. Instead of deploying 49 cameras and allocating a budget of around £75,000, a single drone equipped with a camera can detect defects as small as 10 mm (small size defects) with camera equipment costing approximately £2,000.

To effectively demonstrate the comparison, consider the example of a wing inspection. Initially, it is necessary to determine the specifications of the drone's camera. The assumption is that the drone flies approximately 1 m above the surface under inspection and that the field of view is limited to an area of  $1 \times 1$  m. This operational setup generates images with appropriately sized artefacts suitable for input into a DL algorithm that has been effectively trained to detect surface defects. Using the camera specification algorithm, as described in Section 3, the Lucid(11) camera equipped with lens(3) can illustrate 10 mm defects at 45 pixels in the captured image.

Another important parameter to take into account is the time required for the drone to completely cover and inspect the area. To approximate this duration, consider the example of the wing area of an A320. Each wing covers approximately 61.3 m<sup>2</sup>, and combined, the total area for both wings is 122.6 m<sup>2</sup> [48] (Fig. 13). If the drone is photographing 1 m<sup>2</sup> patch, it needs approximately 61 images. The 122.6 m<sup>2</sup> figure refers to the flat 'shadow' area of the wing. In reality, the upper (and lower) skin of the wing is slightly curved (aerofoil shape), so the total area of the skin is somewhat larger, which can be approximated roughly to 63 m<sup>2</sup>. To thoroughly survey the wing, the drone is generally flown in a



**Figure 13.** The Airbus A320 wings (Source online in Ref. (50)). The area of both wings is 122.6 m<sup>2</sup>.

‘lawnmower’ (boustrophedon) pattern [49]. This involves making straight, parallel flights separated by 1 m (the camera’s swath), followed by a 180° turn at the end of each line to start the next one. It makes sense to fly each pass along the span (wing tip to wing root) because the A320 half-wing span is about 17.05 m, so each sweep is roughly 17 m long. In this case, covering the entire top of an A320 wing at 0.5 m/s, with a 1 m swath and 5 s per 180° turn, requires on the order of 2.5 minutes total. Although this serves as a basic representation of the coverage pattern, it provides a ballpark estimate that requires additional refinement to accurately reflect the realistic time measurement.

#### 4.3.3. Scenario C – ground assets and human monitoring

The monitoring case focusses on tracking moving assets on the ground, including ground support vehicles, personnel, tool cribs, ground power units (GPUs) and tugs. Typically, height assumptions for viewing depth range from 1 to 2 m, resulting in an average distance of 21.5 m from the ceiling. From an overhead perspective, there is a significant difference in the area occupied by a vehicle compared to that of a human (Table 1). This is why the camera specifications are designed separately for the two distinct scenarios. The maximum speed allowed for terrestrial vehicles is 2 m/s, while the typical average walking speed for a person is about 1.27 m/s. In the first case, the chosen camera records 5 cm per frame, while in the second case, it captures 9.4 cm per frame. For ground vehicles, four cameras are necessary (Fig. 14); however, if the targets include humans, the requirement increases to nine cameras (Fig. 15). The price, which includes a 24-port PoE Ethernet Switch, is approximately £8,800 and £17,200, respectively.

#### 4.3.4. Comparative blueprint synthesis

To consolidate the design-to-cost blueprints discussed, Table 6 contrasts the five scenarios of the CMC system with the MoCap and UWB reference systems. The snapshot highlights how the intended application, bill-of-materials footprint, and technology maturity interact with overall budget, providing a concise decision aid for hangar planners.

In the context of localisation, the three systems tackle the same fundamental issue but employ different physical principles and enabling technologies. MoCap systems achieve the highest fidelity in pose estimation, offering static errors within a submillimetre range and orientation accuracy below one degree, with update frequencies exceeding 100 Hz. This precision depends on maintaining simultaneous visibility of markers to infrared cameras. MoCap systems are a relatively expensive infrastructure.

UWB systems achieve centimetre-level accuracy, generally less than 10 cm, operating at frequencies of tens of hertz using ToF ranging between fixed anchors and tags. Although a single tag cannot directly determine orientation, using multiple tags on the object could be used to infer the relative rotations.

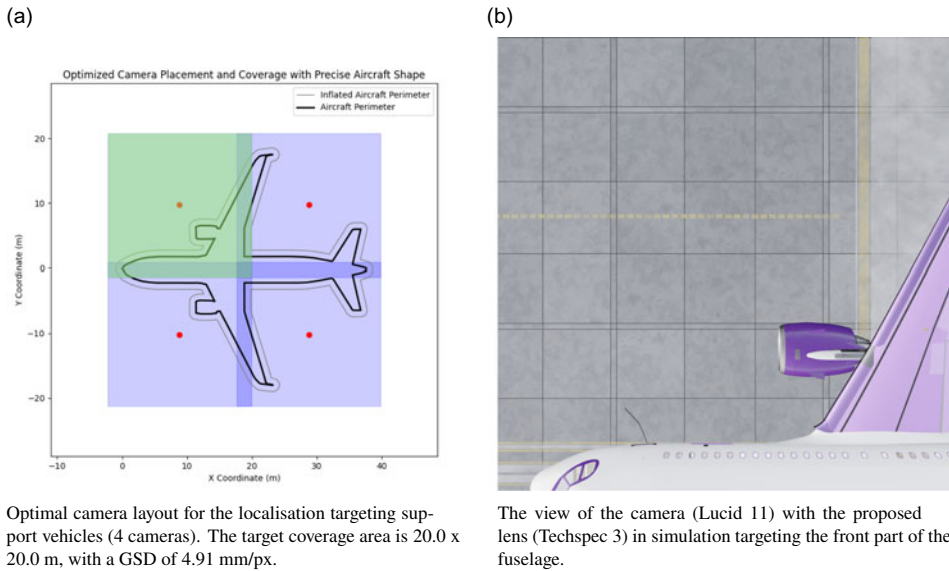


Figure 14. Monitoring scenario targeting ground support vehicles for the camera-based system.

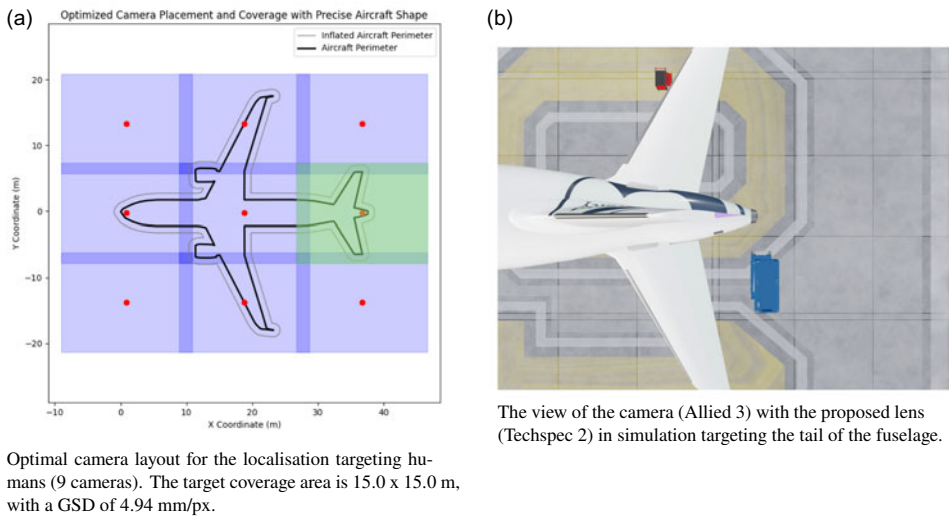


Figure 15. Monitoring scenario including humans for the camera-based system.

UWB remains effective regardless of lighting conditions. Nonetheless, substantial multi-path interference and dense metal obstructions in a hangar can influence range measurements. Economically, UWB technology is positioned at an intermediate cost level.

The CMC system blueprint differs because the target asset does not require a tag or marker. However, models utilised for machine vision reasoning must be trained on representative datasets. One camera grid configuration can establish the exact location of a drone or ground robot within a single bay, another setup can monitor tugs and personnel on the floor and a more intricate and high-resolution layout can inspect the upper surfaces of the aircraft. The system’s performance metrics are influenced by the choice of hardware and the vision model optimisation. There is no commercial solution available for comparison.

**Table 6.** Comparison of localisation and monitoring blueprints proposed

Blueprint	Application	Indicative on-bay equipment	Cost <sup>1</sup> (£k)	Maturity
CMC System	Medium-sized defect detection on upper surfaces	49×20MP rolling shutter GigE cameras, lenses, and 3 PoE Ethernet switches	76.8	Low TRL
CMC System	Drone localisation	15×2.3MP global shutter GigE cameras, lenses, and 1 PoE Ethernet switch	16.5	Low TRL
CMC System	Ground robot localisation	8×7.1MP global shutter GigE cameras, lenses, and 1 PoE Ethernet switch	12.8	Low TRL
CMC System	Vehicle tracking	4×12.3MP global shutter GigE cameras, lenses, and 1 PoE Ethernet switch	8.8	Low TRL
CMC System	Personnel tracking	9×12.4MP global shutter GigE cameras, lenses, and 1 PoE Ethernet switch	17.2	Low TRL
UWB System	Asset localisation / tracking	10–25 anchors, tags, timing units, PoE switches	49.0	High TRL
MoCap System	High-precision localisation	12 IR cameras, active markers, sync hub	190.0	High TRL

**Note:** <sup>1</sup>The price estimate for the CMC system excludes cables, mounts, and computing units.

There are trade-offs to consider between these systems. Firstly, in terms of accuracy and maturity, the MoCap and UWB designs are high-TRL, vendor-supported solutions, exhibiting predictable performance in industrial environments. In contrast, the camera-based system is low-TRL, and its accuracy in localisation, robustness to occlusion and long-term stability in reflective and dynamic hangars remain areas of research. Secondly, regarding temporal behaviour, MoCap regularly surpasses 100 Hz, and UWB tags can operate at several tens of hertz. The proposed camera set-up operates at a few tens of frames per second, which is usually sufficient to track slow-moving objects. Third, in observability, the CMC system estimates both the position and orientation of a rigid body, while UWB requires multiple tags to determine orientations. Lastly, regarding scalability and coverage, the UWB and MoCap systems are restricted to detecting specifically tagged or marked objects. In contrast, once installed, a CMC system can observe everything within its view, having been trained for such targets. This extensive visibility aids in safety and coordination on the ground, but raises issues around data governance and privacy, especially with respect to personnel monitoring.

The comparison of the systems indicates that each one has its own strengths and weaknesses based on the specific requirements. MoCap is the reference for high-accuracy ground-truth and precision robotics localisation. UWB is a mature option for tracking assets and platforms without relying on vision. The CMC system is a flexible system capable of tasks ranging from localising robots and personnel to visually inspecting aircraft surfaces. However, it remains the least developed and relies heavily on strong perception algorithms.

## 5. Conclusions

This study presents the first comprehensive techno-economic analysis for localisation, artefact detection, and monitoring within a smart hangar. Rigorously benchmarking MoCap, UWB and CMC systems across three representative scenarios—robot location, asset tracking, and surface defect detection – enables the demonstration of the interactions between accuracy, coverage and cost within a  $40 \times 50$  m bay. The proposed two-stage optimisation framework for camera selection and placement combines market-driven camera lens selection with optimisation algorithms, demonstrating how a CMC system can be deployed in various operational scenarios. Design-to-cost case studies demonstrate that only 15 global-shutter GigE cameras are sufficient for drone localisation, 9 for monitoring, whereas 49 high-resolution units enable complete airframe midsize defect mapping.

These findings have immediate practical value for MRO planners. The blueprint table furnishes a ready-made bill-of-materials envelope, reducing specification time and derisking procurement. Future work should address two open challenges. First, large-scale certification will require extended endurance trials to characterise failure modes under dense reflections and dynamic occlusions. Second, development and validation of the customised deep-learning approaches to support the camera-based solution back-end. Resolving these issues will propel the transition from isolated demonstrators to fully integrated smart hangars in which sensing infrastructures, robots and human technicians collaborate seamlessly to enhance safety, turnaround time and sustainability.

**Acknowledgments.** We extend our gratitude to Martin Lewis and Andy Ward for their invaluable support and insights on commercially available systems.

## References

- [1] Plastropoulos, A., Zolotas, A. and Avdelidis, N.P. A comprehensive review of robotics-aided aircraft non-destructive inspection towards the smart hangar, *Aeronaut. J.*, 2025, pp 1–32. issn: 0001-9240. doi: [10.1017/AER.2025.10048](https://doi.org/10.1017/AER.2025.10048)
- [2] Masiero, A., et al. A comparison of UWB and motion capture uav indoor positioning, *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci. – ISPRS Archives* 42.2/W13, 2019, pp 1695–1699. issn: 16821750. doi: [10.5194/ISPRS-ARCHIVES-XLII-2-W13-1695-2019](https://doi.org/10.5194/ISPRS-ARCHIVES-XLII-2-W13-1695-2019)
- [3] Van Gerwen, V.J., et al. Indoor drone positioning: accuracy and cost trade-off for sensor fusion, *IEEE Trans. Veh. Technol.*, 2022, **71**, (1), pp 961–974. issn: 19399359. doi: [10.1109/TVT.2021.3129917](https://doi.org/10.1109/TVT.2021.3129917)

- [4] Park, I. and Cho, S. Fusion localization for indoor airplane inspection using visual inertial odometry and ultrasonic RTLS, *Sci. Rep.*, 2023, **13**, (1), pp 1–13. issn: 20452322. doi: [10.1038/s41598-023-43425-y](https://doi.org/10.1038/s41598-023-43425-y)
- [5] Pugliese, R., Konrad, T. and Abel, D. LiDAR-aided relative and absolute localization for automated UAV-based inspection of aircraft fuselages, *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 2021. doi: [10.1109/MFI52462.2021.9591172](https://doi.org/10.1109/MFI52462.2021.9591172)
- [6] Jong, M.C., et al. Handling, inspection and repair of aircraft composites: a pilot study on the awareness of maintenance personnel, *Aeronaut. J.*, 2024, **128**, (1320), pp 211–229. issn: 0001-9240. doi: [10.1017/AER.2023.28](https://doi.org/10.1017/AER.2023.28)
- [7] Moenck, K. et al. Industry 5.0 in aircraft production and MRO: challenges and opportunities, *CEAS Aeronaut. J.*, 2025, pp 1–25. issn: 18695590. doi: [10.1007/s13272-025-00832-3](https://doi.org/10.1007/s13272-025-00832-3)
- [8] Wallis, A.R. and Ruglen, N. On the aerodynamics of “Hangar” type engine test facilities, *J. R. Aeronaut. Soc.*, 1966, **70**, (662), pp 312–320. issn: 0368-3931. doi: [10.1017/S0368393100082183](https://doi.org/10.1017/S0368393100082183)
- [9] Plastropoulos, A., et al. The ‘hangar of the future’ for sustainable aviation, *Aeronaut. J.*, 2024, **128**, (1329), pp 2429–2450. issn: 0001-9240. doi: [10.1017/AER.2024.79](https://doi.org/10.1017/AER.2024.79)
- [10] Ubisense. *Aviation & MRO Asset Tracking – The Ultimate Guide*. <https://ubisense.com/aviation-mro-asset-tracking/>
- [11] Ubisense. Ubisense and MRO Drone launch ‘Smart Hangar’ solution. <https://ubisense.com/ubisense-and-mro-drone-launch-worlds-first-smart-hangar-solution/>
- [12] Aurand, A.M., Dufour, J.S. and Marras, W.S. Accuracy map of an optical motion capture system with 42 or 21 cameras in a large measurement volume, *J. Biomech.*, 2017 **58**, pp 237–240. issn: 0021-9290. doi: [10.1016/J.JBIOMECH.2017.05.006](https://doi.org/10.1016/J.JBIOMECH.2017.05.006)
- [13] Konrad, J. et al. High-accuracy people counting in large spaces using overhead fisheye cameras, *Energy Build.*, 2024, **307**, p 113936. issn: 0378-7788. doi: [10.1016/J.ENBUILD.2024.113936](https://doi.org/10.1016/J.ENBUILD.2024.113936)
- [14] Huang, B. et al. ASD-YOLO: an aircraft surface defects detection method using deformable convolution and attention mechanism, *Measurement*, 2024, **238**, p 115300. issn: 0263-2241. doi: [10.1016/J.MEASUREMENT.2024.115300](https://doi.org/10.1016/J.MEASUREMENT.2024.115300)
- [15] Merriaux, P. et al. A study of vicon system positioning performance, *Sensors*, 2017, **17**, (7), p 1591. issn: 1424-8220. doi: [10.3390/S17071591](https://doi.org/10.3390/S17071591)
- [16] Rahimian, P. and Kearney, J.K.. Optimal camera placement for motion capture systems, *IEEE Trans. Visualization Comput. Graphics*, 2017, **23**, (3), pp 1209–1221. issn: 10772626. doi: [10.1109/TVCG.2016.2637334](https://doi.org/10.1109/TVCG.2016.2637334)
- [17] Meyer, Z. and Pretorius, A. Design of a low-cost optical motion capture system using a multicamera configuration and an asynchronous extended Kalman filter, in du Preez, W., et al. *MATEC Web of Conferences*, vol. **406**, 2024, p 04010. issn: 2261-236X. doi: [10.1051/MATECONF/202440604010](https://doi.org/10.1051/MATECONF/202440604010)
- [18] Lvov, G., et al. Mobile MoCap: Retroreflector localization on-the-go, *IEEE International Conference on Automation Science and Engineering*, August 2023. issn: 21618089. doi: [10.1109/CASE56687.2023.10260562](https://doi.org/10.1109/CASE56687.2023.10260562). **28 The Aeronautical Journal**
- [19] Hansen, J.G., et al. UAV trajectory evaluation in large industrial environments: a cost-effective solution, 2022 European Control Conference, ECC 2022, 2022, pp. 1336–1341. doi: [10.23919/ECC5457.2022.9838352](https://doi.org/10.23919/ECC5457.2022.9838352)
- [20] Zafari, F., Gkelias, A. and Leung, K.K. A survey of indoor localization systems and technologies, *IEEE Commun. Surv. Tutorials*, 2019, **21**, (3), pp 2568–2599. issn: 1553877X. doi: [10.1109/COMST.2019.2911558](https://doi.org/10.1109/COMST.2019.2911558).
- [21] Schroerer, G. A real-time UWB multi-channel indoor positioning system for industrial scenarios, *IPIN 2018 – 9th International Conference on Indoor Positioning and Indoor Navigation*, 2018. doi: [10.1109/IPIN.2018.8533792](https://doi.org/10.1109/IPIN.2018.8533792)
- [22] Fatima, H.S. et al. High precision indoor positioning system for mobile asset management and safety, *Eng. Proc.*, 2022, **20**, (1), p 37. issn: 2673-4591. doi: [10.3390/ENGPROC2022020037](https://doi.org/10.3390/ENGPROC2022020037)
- [23] Karadeniz, G.C., et al. Precise UWB-based localization for aircraft sensor nodes, *AIAA/IEEE Digital Avionics Systems Conference – Proceedings 2020-October*, 2020. issn: 21557209. doi: [10.1109/DASC50938.2020.9256793](https://doi.org/10.1109/DASC50938.2020.9256793)
- [24] Wang, T., et al. Multipath-assisted single-anchor localization via deep variational learning, *IEEE Trans. Wireless Commun.*, 2024, **23**, (8), pp 9113–9128. issn: 15582248. doi: [10.1109/TWC.2024.3359047](https://doi.org/10.1109/TWC.2024.3359047)
- [25] Leugner, S. and Hellbrück, H. Lessons learned: Indoor Ultra-Wideband localization systems for an industrial IoT application, *Proceedings of the 3rd KuVS/GI Expert Talk on Localization*, 2018, pp 2–4. doi: [10.24355/DBBS.084-201807191346-0](https://doi.org/10.24355/DBBS.084-201807191346-0)
- [26] Tatala, O. et al. Quadcopter: design, construction and testing, *Int. J. Res. Eng. Appl. Manag.*, 2018, Special Issue AMET-2018, pp 1–7. doi: [10.18231/2454-9150.2018.1386](https://doi.org/10.18231/2454-9150.2018.1386)
- [27] Husarion. *Panther Robot Specifications*. <https://husarion.com/manuals/panther/overview/>
- [28] Clearpath Robotics. *Husky A300 Robot Specifications*. <https://clearpathrobotics.com/husky-a300-unmanned-ground-vehicle/-robot/>
- [29] Boston Dynamics. *Spot Rorot Specifications*. <https://support.bostondynamics.com/s/article/Spot-Specifications-49916>
- [30] Panero, J. and Zelnik, M. *Human Dimensions and Interior Space: A Source Book of Design Reference Standards*. Whitney Library of Design, 1992, pp 71–74. isbn: 978-0823072712.
- [31] Forklift Safety. *What Are The Recommended Speed Regulations For Forklift Operations?* <https://www.forkliftsafety.com/2023/09/16/what-are-the-recommended-speed-regulations-for-forklift-operations/>
- [32] Transmon Engineering Ltd. *Managing Forklift Speeds*. <https://www.transmon.co.uk/news/managing-forklift-speed-limits/>
- [33] Murtagh, E.M. et al. Outdoor walking speeds of apparently healthy adults: a systematic review and meta-analysis, *Sports Med. (Auckland, N.Z.)*, 2021, **51**, (1), pp 125–141. issn: 1179-2035. doi: [10.1007/S40279-020-01351-3](https://doi.org/10.1007/S40279-020-01351-3)
- [34] Bentham, J. et al. A century of trends in adult human height, *eLife*, 2016, **5**. issn: 2050084X. doi: [10.7554/ELIFE.13410](https://doi.org/10.7554/ELIFE.13410)
- [35] SpanTech. *Narrow Body Aircraft Hangars*. <https://www.airport-suppliers.com/product/narrow-body-aircraft-hangars/>
- [36] Hayduk, R.J. Hail damage to typical aircraft surfaces, *J. Aircraft*, 1973, **10**, (1), pp 52–55. issn: 00218669. doi: [10.2514/3.60196](https://doi.org/10.2514/3.60196)
- [37] Chen, X., Ren, H. and Bil, C. Inspection intervals optimization for aircraft composite structures considering dent damage, *J. Aircraft*, 2014, **51**, (1), pp 303–309. issn: 15333868. doi: [10.2514/1.C032377](https://doi.org/10.2514/1.C032377)

- [38] Moupfouma, F. and Moupfouma, F. Aircraft structure paint thickness and lightning swept stroke damages, *SAE Int. J. Aerospace*, 2013, **6**, (2), pp 392–398. issn: 19463901. doi: [10.4271/2013-01-2135](https://doi.org/10.4271/2013-01-2135). Plastropoulos et al. Techno-economic analysis 29
- [39] Rojas, J.L., et al. Innovative NDT technique based on ferrofluids for detection of surface cracks, *J. Nondestr. Eval.*, 2015, **34**, (4), pp 1–12. issn: 15734862. doi: [10.1007/s10921-015-0309-5](https://doi.org/10.1007/s10921-015-0309-5)
- [40] Fisher, B., Keyser, G. and Deal, P. Lightning Attachment Patterns and Flight Conditions for Storm Hazards ' 80. Tech. rep. National Aeronautics and Space Administration, 1982. <https://ntrs.nasa.gov/api/citations/19830005806/downloads/19830005806.pdf>
- [41] Baaran, J. Study on visual inspection of composite structures, European Aviation Safety Agency, 2009. [https://www.easa.europa.eu/sites/default/files/dfu/EASA\\_REP\\_RESEA\\_2007\\_3.pdf](https://www.easa.europa.eu/sites/default/files/dfu/EASA_REP_RESEA_2007_3.pdf)
- [42] Airbus. *AutoCAD 3-view aircraft drawings*, 2025. <https://aircraft.airbus.com/en/customer-care/fleet-wide-care/airport-operations-and-aircraft-characteristics/autocad-3-view-aircraft-drawings>
- [43] The Boeing Company. *CAD 3-View Drawings for Airport Planning Purposes*, 2025. <https://www.boeing.com/commercial/airports/3-view>
- [44] Mitchell, S., O'Sullivan, M. and Dunning, I. PuLP: a linear programming toolkit for python, 2011. <https://optimization-online.org/wp-content/uploads/2011/09/3178.pdf>
- [45] Forrest, J., et al. Coin-or/Cbc. Tech. rep. doi: [10.5281/ZENODO.13347261](https://doi.org/10.5281/ZENODO.13347261)
- [46] Statista. Aviation industry – aircraft fleet by type. English. Apr. 2025. <https://www.statista.com/statistics/573231/aviation-industry-aircraft-fleet-by-type/>
- [47] Nagymáté, G. and Kiss, R.M. Motion capture system validation with surveying techniques, *Mater. Today: Proc.*, 2018, **5**, (13), pp 26501–26506. issn: 2214-7853. doi: [10.1016/J.MATPR.2018.08.107](https://doi.org/10.1016/J.MATPR.2018.08.107)
- [48] Airbus. Customer Services Technical Data Support and Services. Tech. rep. Blagnac Cedex, June 2024. [https://aircraft.airbus.com/sites/g/files/jlcbita126/files/2025-01/AC\\_A320\\_0624.pdf](https://aircraft.airbus.com/sites/g/files/jlcbita126/files/2025-01/AC_A320_0624.pdf)
- [49] B`ahnemann, R., et al. Revisiting boustrophedon coverage path planning as a generalized traveling salesman problem, *Springer Proceedings in Advanced Robotics*, **16**, 2019, pp 277–290. doi: [10.1007/978-981-15-9460-1\\_20](https://doi.org/10.1007/978-981-15-9460-1_20)
- [50] *Aviation Stack Exchange – What is the total wing surface area of the A320?* <https://aviation.stackexchange.com/questions/54511/what-is-the-total-wing-surface-area-of-thea320>