Computers & Security 154 (2025) 104433

Contents lists available at ScienceDirect

Computers & Security

journal homepage: www.elsevier.com/locate/cose

Full length article

Enhancing maritime cyber situational awareness: A cybersecurity visualisation for non-experts

Dominic Too *, Louise Axon, Ioannis Agrafiotis, Michael Goldsmith, Sadie Creese

Department of Computer Science, University of Oxford, Parks Road, Oxford, OX1 3PR, United Kingdom

ARTICLE INFO

Keywords:
Visualisation
Maritime
Intrusion detection
Cyber–physical systems
Situational awareness

ABSTRACT

Cyber situational awareness is key to mitigating the impacts of cyber threats. However, maritime falls short of its comparative industries, with very little attention given to cyber threats despite the growing concern. In this paper, we explore the use of visualisations as a way to improve the situational awareness of non-experts onboard ships. We designed a visualisation tool with focus on systems that are accessible once onboard. In order to elicit requirements for our visualisations, we conducted semi-structured interviews with experts. We further created a synthetic dataset of attacks that target the systems of ships, which we used to assess the usability of our visualisation. In order to evaluate our visualisations, we conducted a user study with both expert and non-expert users. Our results show that non-expert participants were able to accurately and efficiently detect synthetic attacks targeting ships in an experimental setting, and they were able to use the visualisation to consider what the consequences of these attacks might be. Expert evaluations further suggest the visualisation has merit as a training tool for raising awareness among maritime employees.

1. Introduction

1.1. Background

The Internet of Things (IoT) is a fast-emerging domain, with 13.2 billion IoT connections in 2022, and a forecasted 34.7 billion connections by 2028 (Ericsson, 2023). The Industrial Internet of Things (IIoT), or Industry 4.0, is the convergence of Information Technology (IT) and Operational Technology (OT), to facilitate communication between devices, systems, and sensors, through the application of IoT technologies within industrial contexts. Accenture estimates that the IIoT could add $14.2 trillion to the global economy by 2030 (Accenture, 2015). With the *Global Risks Perception Survey 2022–2023* placing 'Cyberattacks on critical infrastructure' amongst the top risks for 2023 (World Economic Forum, 2023), it is clear that the cybersecurity of IIoT must be given due consideration. Maritime is a sector that falls short of its comparative industries when cybersecurity is considered, in particular when considering the unique OT environments of ships.

Recent legislative efforts in the EU, with the publication of the NIS2 Directive (EU, 2022), have acknowledged the importance of the maritime sector, given that it handles over 80% of the volume of global trade (UNCTAD, 2021). The legislation calls for raising awareness of cybersecurity amongst non-experts and mandates specific security measures in order to increase the maturity posture of all organisations that are active in this sector.

Further efforts are being made globally with focus on technology integrated on ships. The International Association of Classification Societies (IACS) has addressed the need for improved cyber-resilience for both ships and on-board systems with the publication of URs E26 and E27, which came into effect for new ships constructed from January 2024 (IACS, 2023a,b). UR E26 addresses the ship as a whole, considering the integration of systems to maintain cyber-resilience, while UR E27 is focused on the security requirements that individual devices must demonstrate. These unified requirements aim to provide minimum requirements for an effective cyber-risk management system, to improve operational resilience, as well as minimum security capabilities for computer-based systems. These legislative efforts followed the publication of the International Electrotechnical Commission's (IEC) technical requirements in 2021, which aimed at improving the cybersecurity of shipborne radio and navigational equipment (IEC, 2021).

The maritime sector's reliance on increasingly interconnected systems, in particular those onboard, coupled with the critical role it plays in global trade, underscores the need for heightened cybersecurity vigilance. In their 2023 analysis of maritime cyber-incidents, CyberOwl

* Corresponding author.
E-mail addresses: dominic@dominictoo.com (D. Too), louise.axon@cs.ox.ac.uk (L. Axon), ioannis.agrafiotis@cs.ox.ac.uk (I. Agrafiotis), michael.goldsmith@cs.ox.ac.uk (M. Goldsmith), sadie.creese@cs.ox.ac.uk (S. Creese).

https://doi.org/10.1016/j.cose.2025.104433
Received 23 August 2024; Received in revised form 8 February 2025; Accepted 10 March 2025
Available online 21 March 2025

found that over 75% of incidents required crew actions as a response, commonly involving shore-based cyber-practitioners instructing the crew on steps to contain the incident (Kenney and Macdonald, 2023). This is despite 20% of respondents not knowing what actions are required of them during an incident. Furthermore, 33% of survey respondents felt that the biggest challenge in improving cyber-risk management is understanding the level of risk. Cyber situational awareness tools such as visualisations play a pivotal role in this context, where accessible and intuitive tools could bridge the knowledge gap among non-expert personnel, therefore improving the sector's resilience against cyber-risks and aligning with legislative and industry mandates.

To the best of our knowledge, there exist scarce cybersecurity visualisations designed with non-expert users in mind that could provide cybersecurity training and raise awareness of cyber risks (McKenna et al., 2016; Latvala et al., 2017). Furthermore, while some industry-led dashboard products have emerged within the last few years, there exist only limited publications of cybersecurity visualisations that are tailored to systems that ships use and can provide situational awareness to experts dealing with cybersecurity issues. Given the capacity challenge of cyber experts faced globally, a tool useable by non-experts may prove useful, particularly in the maritime industry which not only experiences unique physical isolation challenges, but also suffers from a significant lack of cyber-awareness across the industry.

Our paper contributes to knowledge by: identifying key challenges and design requirements for a visualisation of the cybersecurity of maritime, and ships in particular, through interviews; formulating an attack taxonomy to categorise maritime cyber-attacks, with emphasis on those on vessels; creating a synthetic attack dataset against a ship; designing and developing cybersecurity visualisation to detect cyber-attacks against a ship; providing empirical evidence of the effectiveness of the developed visualisation for non-experts; and conducting expert evaluations of the utility of the visualisation and its applications in maritime.

More specifically, in this paper we provide a visualisation tool that enables non-expert users to identify multiple attacks that affect systems of vessels and allows maritime employees to consider the impact of such attacks. To holistically understand impact, we consider the IIoT connectivity on a modern ship, including that enabled by Programmable Logic Controllers (PLCs).

This paper addresses the following research questions:

RQ1 In what contexts could cybersecurity visualisations be used to enhance the efficacy in monitoring and responding to cyber threats?

RQ2 What are the design criteria and requirements for a cybersecurity visualisation tailored to the detection of cyber-attacks in maritime operations?

RQ3 To what degree can a cybersecurity visualisation be effective for non-expert users to detect attacks?

RQ4 To what extent does a cybersecurity visualisation prove useful amongst stakeholders in the maritime industry?

In what follows, Section 2 provides an overview of the relevant literature in maritime IT and OT systems and cybersecurity visualisations. Section 3 details the methodology used to design, implement and assess the visualisation, while Section 4 provides insights from interviews eliciting requirements for the visualisation. Section 5 provides a novel taxonomy for categorising maritime cyberattacks which is utilised in Section 6 to underpin the visualisation tool, alongside requirements elicited from interviews with experts. Section 7 presents the results of our usability study and Section 8 elaborates on insights from interviewing experts in order to complement our assessment of the visualisation. Finally, Section 9 discusses limitations of our study before Section 10 concludes our paper.

2. Literature review

The state of cybersecurity in maritime

As Industry 4.0 develops, cybersecurity is becoming a pertinent issue; a survey of automation executives by Morgan Stanley cites cybersecurity as the respondents' top concern (Morgan Stanley, 2016). This is particularly evident in maritime environments, where previously 'air-gapping' was a sufficient practical security feature, being isolated at sea, as IT and OT become increasingly interconnected this air-gap disappears and opens these systems up to vulnerability (Byres, 2013). Caponi and Belmont found that maritime has fallen behind comparative industries, such as banking and healthcare, that deal with similar cybersecurity challenges (Caponi and Belmont, 2015).

The number of cyber-attacks is also rising dramatically. The US Coast Guard reported a 111% increase in cyber-incidents from 2020 to 2022 (Department of Homeland Security OIG, 2024). Table A.18 contains notable recent cybersecurity incidents, highlighting the increasing frequency and high impact of cyber-attacks. While many of these are shore-based attacks, there are increasing reports of ship-based attacks, particularly of GNSS and AIS interference. In fact, the real number of cyber-attacks may easily be greater, as many companies refrain from reporting such attacks (Meyer-Larsen and Müller, 2018). This is echoed by Afenyo and Caesar, whose review found that data on cybersecurity incidents is often kept hidden by organisations as a form of 'reputation management' (Afenyo and Caesar, 2023). Mekala et al. state that current security risk prevention mechanisms are 'inadequate' for new threats presented by IIoT, because they are designed for prevention, detection, and response of IoT platforms, rather than IIoT (Mekala et al., 2023). Some of these challenges are shared, as summarised by Yu and Guo, however many are industry-specific (Yu and Guo, 2019).

Bothur et al.'s analysis of security vulnerabilities in a smart ship notes the unique challenges that geographical isolation exposes mariners to Bothur et al. (2017). On top of the traditional issues such as rough weather and pirate attacks, the use of technology, whilst supporting these issues, introduces a host of its own. Some vulnerabilities include:

Automatic Identification System (AIS) AIS is used to broadcast information about a vessel's position. These messages are transmitted unencrypted across unprotected radio systems, which could potentially be used to target ships (International Maritime Organization, 2004). They are vulnerable to Man-in-water Spoofing, disabling of AIS to render ships invisible, triggering fake collision warnings, generating false weather alerts, all of which can trigger wrong decisions being made, and potentially catastrophic outcomes (Balduzzi et al., 2013).

Global Navigation Satelite System (GNSS) GNSS is used by vessels to facilitate geopositioning and enable accurate navigation. Disrupted or manipulated GNSS signals can send ships off course (Humphreys, 2013). Such attacks could lead to collisions, or even violation of international law if ships are navigated on false data (EUSPA, 2023). Jones highlights the growing concern over the impact of GNSS disruption (Jones, 2014).

Industrial Control System (ICS) The ICS connects an array of devices and sensors together, controlling and monitoring temperature, pressure, level, viscosity, flow control, speed, torque, voltage, current, machinery and equipment status (Zaghloul, 2014). In the past, these devices were often 'bolted together' and not designed with security in mind, with data transferred unencrypted (Bothur et al., 2017). More recently, following new regulations introduced by class societies and the IMO, security is increasingly designed into maritime systems, however the extent of implementation varies across the industry. Triton and Stuxnet are examples of malware that specifically targeted ICSs, and caused massive damage to the target organisations (Giles, 2019; Zetter, 2014).

Electronic Chart Display Information System (ECDIS) The NCC Group found the ECDIS vulnerable to attack, with connection to critical systems combined with internet access (Dyryavyy, 2014). It would be possible to subvert sensor data and misrepresent it to the ECDIS; steal or manipulate navigational charts; or compromise the local network and gain access to other data. A compromised ECDIS could lead to loss of life, environmental impact, and financial losses.

Voyage Data Recorder (VDR) The VDR has already seen suspicious data collected, indicating potential tampering to destroy incriminating evidence (Santamarta, 2015). It was also found that some VDRs are vulnerable to buffer overflows, common injection vulnerabilities, and flawed firmware update mechanisms, and that exploitation of vulnerabilities could compromise the confidentiality, integrity and availability of VDR data and services (Hopcraft et al., 2023).

Furthermore, given the interconnectivity of these systems, false information can be propagated by cyber-attacks against other components. For instance, ECDIS charts and routes can be deleted or modified. If a VDR stores that false data, it would provide false information to accident investigators (Söner et al., 2023).

Jones and Tam share a potential high-impact attack against the Port of Valencia, which saw a spear phishing attack insert malware onto a ship's ECDIS via USB vulnerabilities. When triggered by geolocation, the malware sent control messages that ran the vessel aground and physically blocked the port within roughly 2 min and 40 s, with catastrophic consequences to global trade (Jones and Tam, 2024).

The traditionally accepted definition of situational awareness due to Endsley is 'the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future' (Endsley, 1988). Franke and Brynielsson later define cyber situational awareness in line with Endsley, as a subset of situational awareness that concerns the "cyber" environment (Franke and Brynielsson, 2014). Lack of awareness is a sentiment shared across the literature. Focussing on threats that target maritime, Chang et al. identify four challenges: lack of training and experts; outdated systems; risk of being a target; and phishing (Chang et al., 2019). Misas et al. similarly identify low cyber-awareness as one of five main challenges in the sector, as it could have long-term impacts on situational awareness, for example, with the potential overtrusting of technology, and inability to validate information manually (Misas et al., 2022). This is echoed by Heering et al. who found the awareness of cybersecurity in the maritime sector to be 'at a very low level or even non-existent' (Heering et al., 2021). Creese et al. give a unique challenge of maritime to be the lack of a 'security mindset', with local engineers fixing issues on the fly without necessary regard for cybersecurity (Creese et al., 2020). Corallo et al. also conclude that not enough attention is given to the concept of cybersecurity awareness within industrial contexts (Corallo et al., 2022). A study by CyberKeel found a pattern of unawareness of cybersecurity incidents, with cybersecurity being delegated to the IT department rather than involving the C-suite (CyberKeel, 2014). Because of this attitude, there is an unpreparedness across the industry for cyber-attacks. A recent survey by DNV found only 19% of respondents agreeing that their organisation is 'very well prepared' to respond and recover from a cyber-attack at sea (DNV, 2023). Particularly given that the four most common types of breach are related to human error (Klahr et al., 2017), better education and understanding of cybersecurity and secure processes is vital to combat the growing cyber threat.

The International Maritime Organization (IMO) have declared the 'urgent need to raise awareness on cyber-risk threats and vulnerabilities' (International Maritime Organization, 2017). Schinas et al. extend this argument highlighting the need for regulatory requirements. They further coin the term cyber-seaworthiness for ships to denote a list of cybersecurity requirements for ships. They claim that such a concept is required to ensure that the rapid digitalisation that shipping is undergoing, and in particular mega-trends such as marine autonomous surface ships (MASS), will not expand the threat landscape (Schinas and Metzger, 2023). Jacq et al. concur with the need for regulation, as they found that, whilst cyber-monitoring infrastructures are essential to providing quick and appropriate responses, unlike other critical sectors there is no regulation in the maritime sector requiring clear monitoring and detection processes due to the low maturity level to implement this kind of architecture (Jacq et al., 2019). To address the new technology that is adopted in ships, and MASS in particular, Bolbot et al. identify the educational needs and competencies for maritime architects and engineers (Bolbot et al., 2022). Finally, Potamos et al. argue that cyber situational awareness in maritime can be effectively increased by fussing data from vessels, such as radar, AIS, and SIGINT. Transforming such data can be critical in detecting and responding almost real-time to cyber incidents (Potamos et al., 2024).

Cybersecurity visualisations

Visualisation systems help users perform tasks more efficiently, providing a way to easier explore data, such as finding expected or unexpected patterns (Munzner and Maguire, 2015). This cognitive support exploits advantages of human perception, such as parallel visual processing, though Tory and Möller found that the effectiveness of a visualisation is dependent on perception, cognition, and the user's tasks and goals (Tory and Möller, 2004).

There is a large body of literature reporting visualisation tools developed over the past fifteen years to facilitate cybersecurity situational awareness (Jiang et al., 2022). Cybersecurity visualisations have been developed using various different visualisation techniques (e.g., iconic displays, geometrically transformed displays, immersive environments); approaches to allowing users to interact with the visualisation (e.g., filtering, details on demand, linking/brushing); and data sources (e.g., security tools, human input, network traces).

Staheli et al. note that whilst visualisation has emerged as a promising technique to improve effectiveness in an evolving digital threat landscape, current visualisations are too complex or too basic for their intended users, and there is little research on what aspects of a cyber visualisation are effective in supporting the users in operations (Staheli et al., 2014). Furthermore, the authors found that evaluation of these visualisations is lacking, with 46% of reviewed papers having no users at all, and only 10% involving non-expert users. McKenna et al. found that very few tools have considered stakeholders with less technical experience and knowledge (McKenna et al., 2016).

Cybersecurity visualisation for the maritime sector

Zhao and Silverajan present a cybersecurity visualisation platform specifically for maritime, designed to improve multi-stakeholder collaboration (Zhao and Silverajan, 2020). This was aimed to combat the maritime industry's slow response to threats, and poor communication between the different stakeholders, by providing a dynamic interface to effectively coordinate incident responses. Zhao and Silverajan also note that research in cybersecurity visualisation in the maritime domain is limited.

The use of cyber ranges to visualise cybersecurity for mariners, for use in training and awareness-raising, is an evolving area of research (Tam et al., 2021). Antonopoulos et al. developed a decision-support system integrating IDS and SIEM logs, for maritime cybersecurity personnel dealing with cyber-attacks targeting port infrastructures, simulated in a cyber range (Antonopoulos et al., 2022). Palbar Misas et al. used a bridge simulator exercise to test participants' situational awareness in response to a simulated ship's navigational systems being compromised, and identify key situational awareness challenges and training needs (Misas et al., 2024).

The human element of maritime cybersecurity training has also been studied. Erstad et al. demonstrate the use of human-centred design for maritime cyber-resilience training, and argue that "maritime simulators present an effective training solution for new cyber-related incidents" (Erstad et al., 2023). Potamos et al. propose the use of structured learning to develop capacity against the rising threat of ransomware in maritime through awareness training, introducing a cyber range to achieve active learning through the simulation of offensive and defensive actions, and expressing that the engagement of all maritime stakeholders is critical to develop the necessary capabilities (Potamos et al., 2023). Haugli-Sandvik et al. found that deck officers perceive the risk to operational technology as significantly lower than risk to information technology, and suggest that training targeting operational technology might improve the deck officer's comprehension of cyber-consequences and improve their risk perception (Haugli-Sandvik et al., 2024). Oruc et al. present a modular approach to maritime cybersecurity training, to improve awareness across various roles within organisations, from office workers to seafarers (Oruc et al., 2024).

The gap

We have established that very few visualisations exist, in particular focusing on data from ships. Current visualisations, such as from Potamos et al. provide graphical interfaces for cyber range platforms and are extremely useful for running exercises and educating maritime personnel in ransomware (Potamos et al., 2023). However, there is still a gap in visualisations focusing on data from ships to identify attacks and educate stakeholders who are not IT or cybersecurity experts. This is echoed by Latvala et al. who note that the underlying assumption is that visualisations are intended to be used by experts (Latvala et al., 2017). However, IT operations at sea are often assigned to deck officers, who are not necessarily cyber or IT experts (Škrlec et al., 2014). Due to a lack of cyber experts, it is infeasible to have someone permanently onboard, so cyber events are usually escalated to shore cyber expertise centres (Jacq et al., 2019). Since the maritime industry must deal with the unique issue of physical isolation, and often patchy (or easily interrupted) connections to shore, it would be beneficial to have a visualisation that can be used by non-expert mariners. Moreover, Potamos et al. have provided evidence that fussing data from ships (e.g. radar, IES) will enhance maritime cyber situational awareness (Potamos et al., 2024). In addition, a non-expert-friendly visualisation of data from ships could help combat the current lack of cyber-awareness and education that is prevalent in the maritime industry, and raise the level of situational awareness across the board. By combining the promising potential of visualisation as a training tool with the target audience of mariners, who have been identified as key, yet unprepared, in cyber-incidence response, our research explores the effectiveness of a targeted visualisation for detecting cyber-attacks and improving cyber situational awareness.

3. Methodology

In order to answer our research questions, we adopted methodological approaches from various disciplines. Our systematic literature review was followed by interviews with a number of maritime experts in order to elicit requirements for the visualisations. Once the interview data was analysed using thematic analysis, we proceeded with creating a taxonomy based on publicly available data on threats that ships experience. We designed the visualisations to reflect our findings from the interviews and the attack taxonomy and created a synthetic dataset to evaluate them. Finally, we conducted a usability experiment with lay users to understand the effectiveness of the visualisations based on attacks synthesised in our dataset before concluding our research by interviewing experts to identify context of use. Fig. 1 illustrates the methodology we followed in this paper, indicating the research questions addressed at each stage.

3.1. Interviews

By conducting interviews, we intend to capture the current state of cybersecurity in the maritime industry, determine the uses of visualisations in maritime, and identify any areas that might be overlooked, particularly from the perspective of those active in the field. This helps us to answer research questions **RQ1** and **RQ2**.

Following guidance due to Harrell and Bradley, we adopt a semi-structured interview approach (Harrell and Bradley, 2009), where we start with a list of questions for every interviewee but we allow the participants to bring into the conversation new concepts based on their experience. This enabled us to adapt our questioning based on the flow of conversation, leading to a more nuanced understanding of participants' perspectives. This flexibility ensures that essential topics are covered, whilst also allowing room for emergent topics to be examined, therefore increasing the relevance and comprehensiveness of the data collected.

We use a judgemental sampling method, finding known experts and professionals working in our desired field, alongside some snowball sampling. However, we take care to avoid the overrepresentation of a single, networked group.

We utilise three main types of questions in our interviews: A grand tour, which reveals the participant's background, allowing us to understand where their opinions are formed, and what perspective they may offer; experience questions enable us to draw directly from participants' experiences, with specific examples of what they have encountered whilst working in the subject area; and cover term questions allow us to distinguish between groups that we have perceived from our prior research, and perhaps give more context to the different objectives that exist in different domains.

We develop an interview protocol, defining the general structure of our interviews. This compels us to carefully consider and clarify the information desired from the interviews, and reduces the effect of interviewer bias by ensuring consistency across interviews, improving comparability between responses from different participants. Furthermore, by using probes we reduce the risk of social desirability bias and ensure reliability of the data, by enabling us to clarify ambiguous responses. Probes also allow us to remain broad in our questions, whilst not losing focus of the aim of our interviews, and can allow for unexpected data to emerge (Jacob and Furgerson, 2012). The protocol is given in Appendix B.

We leave this protocol open to on-the-spot revisions, allowing for emergent design depending on the responses, especially if unexpected (Jacob and Furgerson, 2012).

After developing this protocol, as suggested by Barriball and While, the draft was reviewed to assess the quality and validity of the questions proposed, ensuring that the interview is complete in purpose and appropriate for the participants (Barriball and While, 1994). This stage is important to verify that no leading questions or ambiguous language is used, which might otherwise elicit false responses and skew the results. We finally analyse the data using thematic analysis (Braun and Clarke, 2006).

3.2. Effectiveness study

Banissi et al. highlight the importance of usability evaluations, a lack of which can result in potentially useful visualisations not being accepted or expanded on due to lack of evidence to encourage adoption, and convincing but less useful ideas being promoted (Banissi et al., 2014). Similarly, our literature review found a lack of user evaluations of visualisations particularly for non-experts. Therefore, we present empirical evidence of the effectiveness of our visualisation through user studies, assessing the accuracy and efficiency of attack detection, based on key concepts provided by Ware (2008) for designing visualisations. This contributes to answering research question **RQ3**.

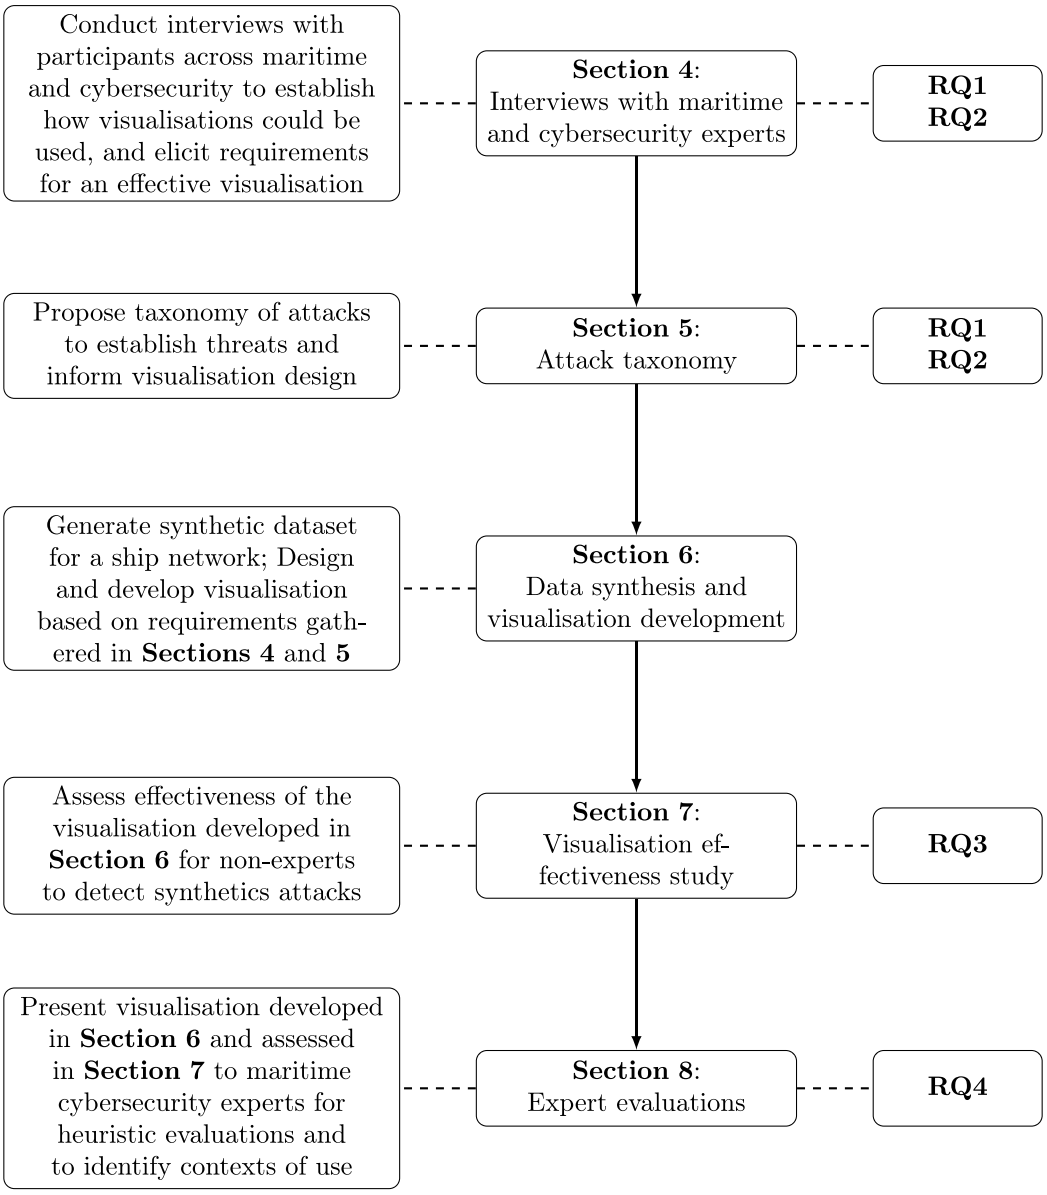

Fig. 1. Overview of research methodology.

We follow a similar methodology to Axon et al. for our effectiveness study (Axon et al., 2019). The study was carried out with students from a range of subject areas. By extending the study to participants both with cybersecurity knowledge, and those with no technical experience, we are able to compare the effect of this knowledge on the efficacy of our visualisation, and offer an analogy to the intended, non-expert users. We also ran a pilot study, to ensure the study was feasible and identify any immediate weaknesses or design issues.

3.2.1. *Pre-training task*

Participants were presented with the visualisation under a pseudo-random attack scenario, with no prior explanation of the visualisation or attack. They were asked to click the mouse on areas of the visualisation that they might think were abnormal or indicative of an attack. After the scenario was finished, they were asked to explain why they clicked, and to describe what they think may have happened in the scenario.

The aim of this task is to assess the intuitiveness of the visualisation, and to compare the effectiveness of the visualisation pre- and post-training.

3.2.2. *Training*

In training, the participants were introduced properly to the visualisation, with the mappings described for each view. This was exemplified under a 'normal' scenario, with no attacks occurring. This scenario was repeated at the participant's request, pausing for explanations of each component and allowing any questions.

3.2.3. *Post-training task*

Each participant was shown all attack scenarios, in a pseudorandomised order to avoid anchoring bias. As in the pre-training task, the participants were asked to click when, and where, they identify abnormal behaviour. The timing and coordinates of the mouse clicks were recorded for quantitative analysis. They were encouraged to vocalise their thoughts throughout the scenario playback. For the scenario used in the pre-training task, we order this last for post-training, in order to directly compare the effect of training whilst minimising memory bias.

After all of the scenarios were completed, the participants were asked for qualitative feedback of the visualisation.

3.2.4. Measuring performance

For each scenario, the application recorded the time at which participants made mouse clicks, relative to the scenario playback, as well as the coordinates of the click. We can thus ascertain the point at which the participant 'detects' an attack. Therefore, we measure detection accuracy and efficiency as follows:

Accuracy

- click in the attack time window AND click in the correct location → true-positive
- click outside the time window OR click in the wrong location → false-positive
- no click in the time window → false-negative

Efficiency

- We calculate the time difference between the beginning of the attack, and the participant's first true-positive click.

We also take into consideration the participants' verbal communication throughout, where they may explain incorrect timings or locations of clicks.

3.3. Evaluation

Whilst our user study captured the demographic of non-expert users, it did not consider the primary intended users' maritime knowledge. Therefore, we conducted expert interviews with participants in the maritime industry to evaluate the suitability of our visualisation. This contributes to answering research question **RQ4**.

We use heuristic evaluations to analyse the visualisation's usability, following advice due to Nielsen and Molich to have three to five evaluators (Nielsen and Molich, 1990). We aim to receive general feedback on the visualisation from experts, and determine how it might fit into the current workflows. The interview protocol is given in Appendix E.

4. Interviews

4.1. Qualitative analysis of interviews

Our interview recruitment reached 6 participants, of both academic and industry backgrounds, at 5 different organisations. 3 participants had primarily cybersecurity backgrounds, whereas the remaining had maritime backgrounds. Participant A1 is an academic researcher specialising in maritime cybersecurity, with expertise spanning ship and port security, autonomous vessels, and offshore structures. Participants B2 and B3 are both executives at a cybersecurity consultancy, which works closely with maritime organisations to provide emergency cyber-response, and cyber-training across the maritime ecosystem. B3 holds a position as an elected board member of an influential maritime organisation. Participant C4 is an underwriter with over 20 years in the maritime industry, including roles in ship management, surveying, and loss prevention, and originally a navigator by trade. Participant D5 is a debt recovery expert and director of an independent marine recovery agency, with more than 20 years of experience in maritime claims handling. Participant E6 is a deck officer on cruise ships, and currently studying towards chief mate certification. The demographic of participants is found in Table 1. One organisation opted to have 2 participants in a single interview (B2 and B3). All interviews took place over live video calls.

After transcribing the interview recordings, or referring to notes where recordings were not consented to, we conducted thematic analysis of the resulting qualitative data (Braun and Clarke, 2006). We opted for a blended approach of inductive and deductive coding, first defining a set of a priori themes that we expect to appear based on our literature review. For each interview, we identified any emergent themes, capturing complexities that may have been overlooked without relying on assumptions or biases. The codebook can be found in Appendix C.

Table 1
Interview participant demographics.

ID	Job Title	Sector	Cyber-Skilled
A1	Researcher	Academia	✓
B2	Consultant	Cybersecurity	✓
B3	Consultant	Cybersecurity	✓
C4	Underwriter	Marine Insurance	
D5	Collector	Marine Insurance	
E6	Deck Officer	Passenger Ships	

Key challenges

Our interviews revealed a range of key challenges relating to cybersecurity.

Forensics were an issue raised by multiple participants (A1, B2, B3, C4), where there was insufficient data following incidences to attribute them to cyber-attacks, leaving us to speculate. This also meant that there was a lack of understanding or improvement following any incidents, as the forensics are not there to determine the root cause. A1 remarked: *"We don't actually learn our lesson [...], we just say, 'Oh, no, there was a cyber-attack', maybe increased protection a bit, but no good understanding of what actually happened and how we can mitigate that"*. Furthermore, this often results in the attribution of attacks to either human error, where in truth the correct decision was made with the (wrong) data, or mechanical failure.

There is a notable lack of collaboration in the maritime industry, and reporting of cybersecurity incidences (B2, B3, D5). In part, this is due to reputation management, similar to findings in our literature review. D5 explained: *"It does not look good, when your system has been breached, or you have made a mistake. I think there is still a taboo around this [...] some people will be scared to say anything"*. B3 also noted that, compared to other industries such as banks or water companies, there is no [default] central reporting portal for organisations to share and monitor current maritime cyber-attacks.

Policy was also a recurring issue (A1, B2, B3, C4, E6). B2 explained that there are no regulators that impose fines, and even after IMO guidance updated in line with cyber-focussed best practices, the industry did not see the change they expected. Rather, they suspect that insurance policies are key to driving change. However, as C4 explained, in its current state, insurance does not consider cybersecurity as a separate threat: *"[Cybersecurity] has practically zero impacts to an underwriter. [...] For example, if the cyber-attack affects the rudder of a ship, and causes the ship to collide into another ship, it's not called a cyber incident, it's called a collision. So the head of cover [...] would come under the collision, not cyber-attack"*. Furthermore, an issue arises from the jurisdiction in which vessels are registered. Since flag states and classification societies can impose regulations on the standards of vessels, they have the influence to promote good cyber behaviour. However, shipping companies are able to register with different flag states, applying the cheapest areas to operate in and around, so such regulation is easily circumvented. Similarly, classification societies have a *'commercial pressure'* (C4) to continue in line with the rest of the industry — i.e. if they impose some higher level of standard, ship owners can choose to register their vessels to competing classification societies.

Situational awareness emerged as a pivotal challenge (A1, B2, B3, C4). B2 and B3 noted that there is a lack of understanding of anything to do with cyber; particularly with crew coming from all over the world, more frequently now with their own (unsecure) devices, this low-level of cyber-awareness can easily result in malware introduced by the crew. In addition, as mentioned previously, this lack of awareness often leads to ignorance of attacks, assuming either human error or mechanical failure. C4 explained: *"An engineer will be called and go, 'why is that valve closing or opening?' [...] He'll think it's a malfunction, he'll go fix it"*.

A number of specific attacks were also mentioned, namely malware (A1, B2, B3, C4, E6), phishing (B2, B3, C4), and ransomware (B2, B3, C4).

Training

Almost all participants indicated some lack of training in maritime (B2, B3, C4, D5, E6). B3 explained that skilling crew members is not mandated, and training remains internal. E6 described their own training for chief mates certificate, *"not particularly"* touching on anything cybersecurity-related, although it may have been mentioned in relation to security duties, *"but I think that there isn't a specific cybersecurity or cyber element to training as a deck officer"*. Similarly D5, when undertaking an LLM in marine insurance, had *'zero'* cyber aspect.

However, there is some basic training that seems to be industry standard (A1, B2, B3, E6). Generally, this training takes the form of eLearning videos covering *"high-level hygiene"* (A1). E6 described their training: *"a series of videos [ending in] a multiple choice exam. It's very basic, but it gives us a decent understanding not to click on emails sent by the King of Africa or something like that"*. This specific company trained only manager and officer roles, with the intention that it will *"filter down through them"*.

By contrast, there are some effective cybersecurity training solutions. Immersive ship simulators (A1), have shown that cadets undertaking this training are *"much more aware [that] this might look like mechanical failure, but it could also be a cyber-attack"*. However, whilst successful, this is a unique form of training that does not exist in many places, and cost implications mean it is unlikely to be implemented throughout the industry.

Visualisations

Two participants did not indicate any visualisations that might be used (C4, D5), with C4 explaining: *"there was no visualisation onboard ships themselves"*. However, there are some basic visualisations in use (A1, B2, B3). In some cases, these visualisations are *"basic video clips"* (A1), displaying a cyber-attack taking place. These are aimed across the board at crews and executives, so can suffer from different audiences having different focuses. More generally, as described by B2 and B3, visualisations take the form of KPIs and dashboards measuring the maturity of certain projects or policies, aimed at an executive level.

Participants also identified limitations to some visualisations. As mentioned previously, different stakeholders have different interests. A1 explained: *"Sometimes they need the big flashy thing to pay attention. [Others] care more about the small details in life, if the power goes out for two hours, or the water, Netflix"*. In addition, accessibility is a concern. B2 explained that with mixed crews arriving from all over the world, visualisations that are language-dependent cannot be understood by everyone onboard.

5. A taxonomy of attacks and how to detect them

5.1. Attack taxonomy

According to Hansman and Hunt, the purpose of an attack taxonomy is to establish a standardised means of categorising cyber-attacks, whilst providing a holistic approach to classifying attacks (Hansman and Hunt, 2005). A taxonomy can aid us in establishing the threat landscape, and therefore help us to answer research question **RQ1**.

Wu and Moon propose a taxonomy for CyberManufacturing System attacks over 4 dimensions: vector; impact; target; and consequence, similar to the dimensions proposed by Hansman and Hunt, but each dimension is subdivided into 'cyber' and 'physical' (Wu and Moon, 2017). Other IoT taxonomies have been proposed, such as Sasi et al. who divide attacks based on domain (Sasi et al., 2023), or Krishna et al. who divide attacks based on the architecture layer in which they occur (Krishna et al., 2021). We choose to follow the same dimensions suggested by Wu and Moon as they are effective for multi-stakeholder industries such as maritime, where the exact architectural location or domain of the attack are not as great a concern as the cause and effects. Furthermore, given the need for raising cyber-awareness indicated by our interviews, the categorisation of impacts and consequences can further emphasise the necessity of cybersecurity to stakeholders with less or no technical knowledge. We modify the taxonomy to apply not just to CyberManufacturing, in Fig. 2.

5.2. Attack analysis

Let us consider an example attack analysis using the proposed taxonomy. We choose to consider the real-world incident faced by the Port of Antwerp in 2011 (Direnzo et al., 2016). Fig. 3 illustrates this incident, divided into the following steps:

- Door-opener — Spear phishing emails targeting port authorities and shipping companies introduce malware into network.
- Stage 1 — Once the network was infected, hackers were able to obtain confidential information regarding shipments of specific containers, modify shipments, and obtain security codes to access containers.
- Door-opener 2 — After being discovered and prevented from further infections, attackers broke into the facility and installed keyloggers on computers.
- Stage 2 — With data obtained through the keyloggers, they were able to continue their operations.
- Payoff — Attackers gain physical access to the port with stolen passcodes and information regarding location and times of deliveries, smuggling drugs and weapons through the containers.

We can therefore divide this incident into 4 attacks, which we classify individually with our taxonomy, as illustrated in Table 2. By breaking down attacks in this way, we gain a deeper understanding of exactly how the attacks work and the impact that they have, which can be used to design a visualisation capable of signifying such attacks, or their consequences, to the user.

5.3. Detecting and preventing attacks

By considering the categories defined in our taxonomy, we examine how we might detect or prevent such attacks, including appropriate articles exemplifying them. This can then be used to inform the design of our visualisation (**RQ2**).

Malware This covers a range of software: ransomware, spyware, viruses, worms, amongst others. Ransomware in particular is rife in maritime; a 2023 survey by CyberOwl found nearly 14% of respondents have paid a ransom (Kenney and Macdonald, 2023). Malware detection is a long-standing challenge in cybersecurity, with various approaches but no catchall method (Aslan and Samet, 2020). However, we can often detect malware by its impact and consequence, which we will explore through visualisation.

Phishing, Social Engineering, Impersonation These attacks exemplify the human aspect of cybersecurity. Often used as door-openers, they require some interaction, physical or digital, with a victim, to acquire information that compromises a system. In a 2021 report by the ICS, phishing attempts were the most common attack vector by threat actors (ICS, 2021). However, each of these attacks can be detected and prevented by basic cybersecurity awareness, as highlighted in our interviews and recognised as urgent by the IMO in 2017 (International Maritime Organization, 2017).

USB Devices Insecure USB ports provide an entryway for malware to be installed even onto segmented OT systems; maritime has seen several cases of air-gapped networks being infected through compromised USB devices (Cimpanu, 2018). This problem is particularly faced by the ECDIS, which historically did not have anti-virus, and is often updated via USB (Baraniuk, 2017). Modern ECDIS do typically have requirements for up-to-date anti-virus, however they are not infallible and USB can still be an attack vector for installing malware. In their 2021 report, the ICS included multiple incidents of malware being introduced

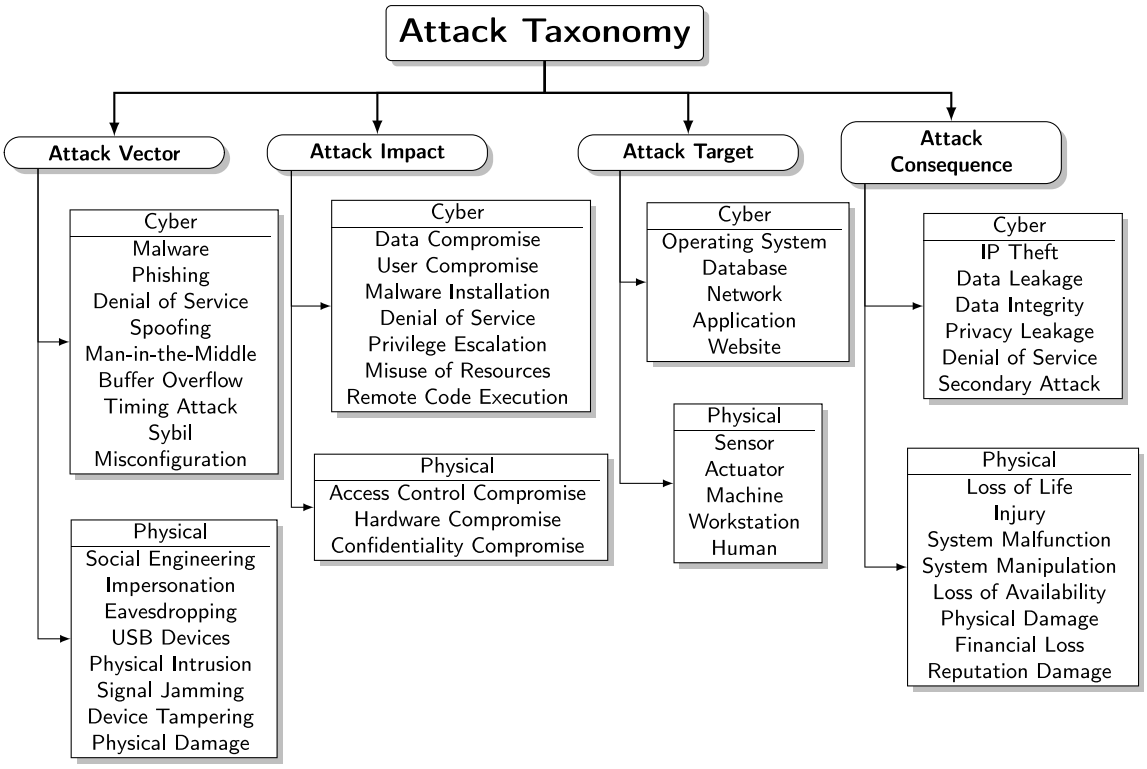

Fig. 2. Attack taxonomy.

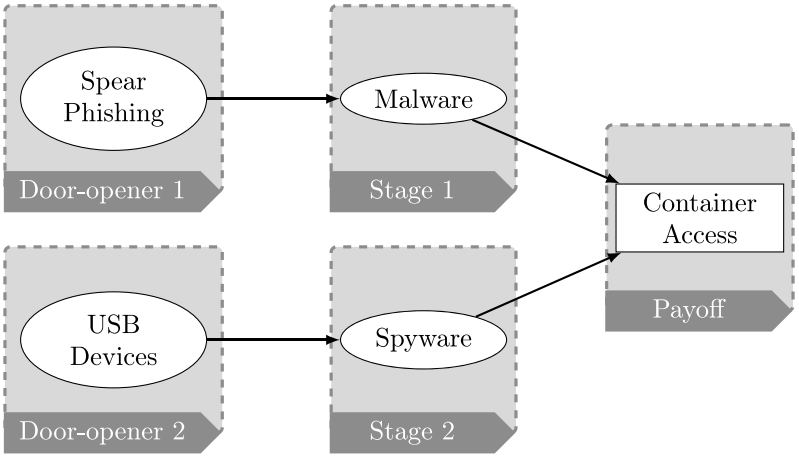

Fig. 3. Port of antwerp attack.

Table 2
Analysis of Port of Antwerp Attack.

Attack	Attack Vector	Attack Impact	Attack Target	Attack Consequence
Door-opener 1	Cyber: Phishing (Spear Phishing)	Cyber: Malware Installation	Physical: Human	Cyber: Secondary Attack
Stage 1	Cyber: Malware	Cyber: Data Compromise	Cyber: Database	Cyber: Data Leakage
Door-opener 2	Physical: USB Devices	Cyber: Malware Installation (Spyware)	Physical: Workstation	Cyber: Secondary Attack
Stage 2	Cyber: Malware (Spyware)	Cyber: User Compromise	Cyber: Application	Cyber: Data Leakage

to vessels through active USB ports (ICS, 2021). Prevention of such attacks is easily solved, by having strict policies on the connection and trusting of external devices, as well as basic cyber-awareness training for crews.

Denial of Service Denial of Service (DoS) attacks flood the network with excessive traffic, to slow or stop genuine users from accessing resources. Several ports, such as the ports of Amsterdam and Groningen in 2023, have reported incidents of DoS Attacks (Jacob et al., 2023). We can detect a DoS attack by analysing the network traffic. This makes DoS one of the easier attacks to visualise, and there are several examples of this available (Kalwar et al., 2020; McAfee, 2013).

Spoofing Spoofing attacks see attackers masquerade as a trusted source, feeding false information to the victim. There are several forms of spoofing attacks: IP spoofing, where the attacker falsifies the source IP address; or, more pertinent to maritime, GNSS spoofing, which has caused ships to alter their course (Humphreys, 2013; EUSPA, 2023; Goward, 2017). These are somewhat difficult to detect. The Galileo OSNMA offers authenticated navigation messages that protect against GNSS spoofing (EUSPA, 2021). However, this requires a GNSS receiver that can leverage the service — shipping companies, as noted in our interviews, are often reluctant to invest in ship upgrades. We will explore ways to use visualisations to detect GNSS spoofing.

Man-in-the-Middle An attacker can insert themselves into communications between two parties, potentially eavesdropping or impersonating either party. The AIS has been identified as vulnerable to MitM attacks, by tampering or replaying AIS communications (Storm, 2013). These attacks are notoriously difficult to detect. There are some methods, such as inspecting packets to determine latency, or by noticing disruption of a service (Mohanakrishnan, 2022). However, the best method to prevent is by using secure connections — which is not an option for the AIS. We can attempt to use visualisations, particularly for the AIS, to detect suspicious activity.

Buffer Overflow If an attacker is able to write outside the bounds of allocated memory, it can cause data to be corrupted, the application to crash, or cause the execution of malicious code (OWASP, 2024). The vulnerability has been shown to be possible on common maritime communications equipment, causing remote code execution (CISA, 2015). Prevention is typically through the use of safe buffer handling functions, non-executable stacks, and address space layout randomisation (Lightner, 2024).

Timing Attack This form of attack exploits the processing time of functions to obtain some knowledge of a system. The AIS has been identified as being vulnerable to timing attacks, by delaying the transmissions, repeatedly preventing communications and thus allowing vessels to 'disappear' (Balduzzi et al., 2014). Timing attacks are generally prevented by ensuring constant time of execution for sensitive operations, or adding some random delay to prevent any information being derived from timing.

Sybil A Sybil attack affects peer-to-peer networks, where a single malicious entity operates under multiple fake identities, with the aim of undermining the authority by gaining a majority influence in the network (Imperva, 2024). Whilst there have not been instances of Sybil attacks in maritime, Rabieh et al. gave example scenarios for autonomous vehicles, describing Sybil attacks fabricating road accident reports, or traffic congestion causing the vehicle to change route (Rabieh et al., 2015). We can map this to a similar scenario for maritime, for example, simulating congestion on popular trade routes by having Sybil ships send false AIS messages to a base station, with the intention of making ships choose to take a different route.

Misconfiguration If particular settings or privileges are improperly configured, an attacker could take advantage of this flaw to gain unauthorised access. We can extend this to include firewalls, open ports, and similar vulnerabilities that require security configurations. Prevention of misconfiguration attacks largely falls into following best practices to ensure proper configuration and setting of default settings (OWASP, 2021).

Physical Intrusion, Physical Damage, Device Tampering These kinds of attacks generally rely on physical security for prevention. Under the maritime context, physical intrusion could cover pirate attacks and intrusive boarding of ships. There are several commercial solutions available to discourage attackers (Maritime Security Alliance, 2024). Patino and Ferryman present a threat detection method based on the trajectory of nearby ships (Patino and Ferryman, 2016). We can consider how we might develop a visualisation capturing such behaviour.

Signal Jamming This attack involves the prevention of wireless signals, to disrupt communications. Jamming attacks have been recognised as a major vulnerability of the GNSS, with potentially critical consequences from depriving them of their navigation guidance (Medina et al., 2019). This is due to the extremely low power of the signals when they are received. Furthermore, in 2019 the Mexican government found that commercial signal jammers were used in 85% of 3400 reported cargo thefts (Mukherji and Chandele, 2024). Because of the nature of the signals, preventing jamming is difficult, so the first line of defence is detection. Interference (both jamming and spoofing) could be detected by methods such as measuring signal strength, jumps in values, or Doppler shifts. The resulting loss of connection might lend itself well to visualisation, which we will explore in our design.

System Manipulation Once an attacker has infiltrated a system, they may be able to gain access to operational technology through control systems. This is a particular issue in maritime, where we have seen examples of actors remotely controlling ships — the cyber equivalent of a hijacking (Blake, 2017; Demchak and Thomas, 2021). This kind of takeover can go unnoticed until the physical consequences have already taken place (e.g. steered off course) (Tvergrov, 2023). Even after this, as noted in our interviews, it could be mistaken as mechanical failure rather than a malicious attack, due to the lack of forensics and cyber-awareness. We will consider a way to incorporate detection of this into our visualisation.

6. Visualisation design

Intrusion Detection Systems (IDS) are generally categorised into signature-based and anomaly-based detection (Scarfone and Mell, 2007). The former relies on pattern recognition of known attack behaviour, whereas the latter monitors for some deviation from the normal characterisation. We aim towards anomaly-detection based on the signatures of attacks, such that our visualisation does not require prior knowledge of attack behaviour. Furthermore, we postulate that it is easier for a human to notice deviations from normal behaviour, rather than retaining and recalling 'bad' behaviours. However, we must be aware that, as is common with anomaly-based IDSs, dynamic environments can cause a significant number of false positives.

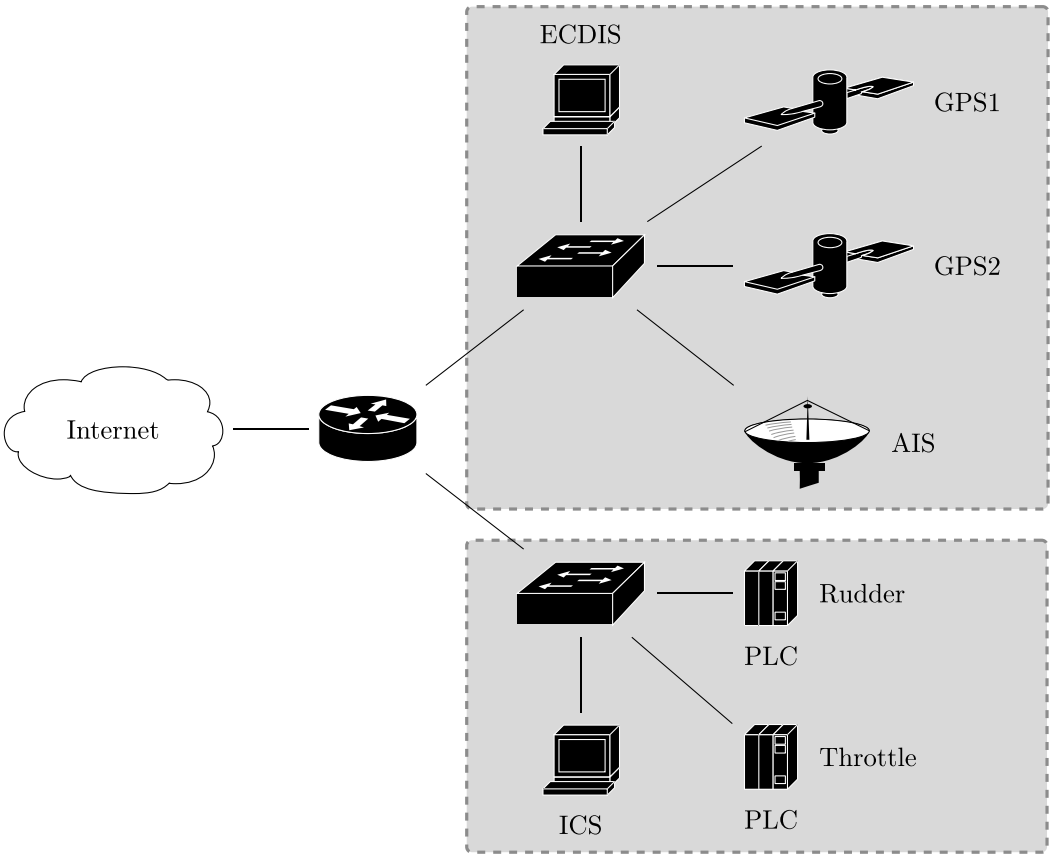

Fig. 4. Proposed ship architecture.

6.1. Network architecture

Potamos et al. and Spravil et al. present maritime architectures, based on the typical command-and-control systems on a ship (Potamos et al., 2023; Spravil et al., 2023). Following these, we propose the network architecture depicted in Fig. 4 to use for our visualisation. This architecture captures the key components of a modern ship with a typical network topology. We will not consider a ship's connection to a shore-based operations centre as it falls beyond the scope of this paper.

Our test network consists of two subnets. The first subnet contains the navigation systems, with 2 GPS receivers and an AIS receiver, each of which send data to the ECDIS. Communication between the receivers and the terminal is done through the UDP protocol. The second subnet contains the industrial control systems, which in this case we have simplified to two PLCs controlling a rudder and throttle. Communication between the PLCs and the terminal is done through the Modbus/TCP protocol.

6.2. Data synthesis

Finding real data is difficult. As noted in interviews, the lack of collaboration in the maritime industry means that it is hard to find real examples of network data. As such, we synthesise our own. We create a method to generate synthetic packet captures of our network in Fig. 4, with realistic traffic between components. We then craft a dataset of network captures under different attack scenarios.

We devise Algorithm 1 to simulate the movement of a ship, given some input rudder and throttle controls. This generates a ship trajectory along with rudder and throttle data, so we can accurately mock responses to data requests. At each point in time, we bring the rudder angle closer to the target defined in the input, by an arbitrary function determining how fast the rudder turns. For simplicity, we ignore external factors such as the hydrodynamic forces and resistance, instead modelling the rudder as turning at a constant rate. Similarly, we bring the engine RPM closer to the target based on the throttle power defined in the input. From these values, we can calculate the ship's trajectory.

Calculating the turning radius is non-trivial, requiring consideration of external factors such as depth, draught, cargo distribution, weather conditions, among others (United States Naval Academy, 2024). Typically, commercial ships have manoeuvring characteristics determined by sea trials, which are catalogued and the turning radius is predicted based upon this. Zelazny presents an approximation of the force acting on the rudder (Zelazny, 2014), based on the calculations described by Inoue et al. (1981). However, these approximations are only evaluated on large bulk carriers and container ships, and do not appropriately reflect the forces on our much smaller simulated ship. We instead approximate the ship's course by a quadratic function on the angle of the rudder. Whilst not necessarily accurate, it is a good enough approximation for the demonstration of our visualisation.

To generate AIS data, we observe real-time traffic in our scenario locations using MarineTraffic.[1] From this, we choose an appropriate number of ships to generate signals for on a basic trajectory replicating observed behaviour. We limit these signals to within 20 nautical miles of our ship, a typical range for commercial AIS transponders (ICOM UK, 2024).

In this way, we can generate realistic network traffic between the ECDIS components (AIS and GNSS), and the ICS components (Rudder and Throttle PLCs), by crafting UDP and Modbus/TCP packets respectively with Scapy.[2]

[1] https://www.marinetraffic.com/
[2] https://scapy.net/

Algorithm 1 Generate Ship Data

Require: Input $RudderWrites$ and $ThrottleWrites$
  $t := startTime$
  Initialise $RudderTarget$, $RudderAngle$, $Throttle$, RPM, $speed$, $course$
  while $t < endTime$ **do**
    if $\exists write \in RudderWrites$ *s.t.* $write.time = t$ **then**
      $RudderTarget \leftarrow write.val$
    end if
    if $\exists write \in ThrottleWrites$ *s.t.* $write.time = t$ **then**
      $Throttle \leftarrow write.val$
      $TargetRPM \leftarrow func(write.val)$
    end if
    if $RPM \neq TargetRPM$ **then**
      $RPM \mathrel{+}= func(RPM, \Delta t)$
    end if
    if $RudderAngle \neq RudderTarget$ **then**
      $RudderAngle \mathrel{+}= func(RudderAngle, \Delta t)$
    end if
    $speed = func(RPM, speed, \Delta t)$
    $course = func(RudderAngle, course, \Delta t)$
    $latitude \mathrel{+}= speed \cdot \cos(course)/111111.1$
    $longitude \mathrel{+}= speed \cdot \sin(course)/(111111.1 \cdot \cos(latitude))$
    yield $(t, latitude, longitude, speed, course, rudderAngle, Throttle, RPM)$
    $t \mathrel{+}= 1$
  end while

Table 3
GPRMC Structure.

Field	Structure	Description
1	$GPRMC	Log header
2	utc	UTC of position (hhmmss.sss)
3	pos status	Position status (A = valid, V = invalid)
4	lat	Latitude (DDmm.mm)
5	lat dir	Latitude direction (N = North, S = South)
6	lon	Longitude (DDDmm.mm)
7	lon dir	Longitude direction (E = East, W = West)
8	speed	Speed over ground, knots
9	course	Course over ground, degrees true
10	date	Date (ddmmyy)
11	mag var	Magnetic Variation, degrees
12	var dir	Magnetic variation direction (E = East, W = West)
13	mode	Positioning mode indicator
14	*xx	Checksum

Table 4
AIVDM Structure.

Field	Structure	Description		
1	!AIVDM	Log header		
2	fragments	Total number of message fragments		
3	fragment	The fragment count		
4	sequence num	Sequence message ID		
5	radio chan	Radio channel code		
		Field	Structure	Description
		1	type	Message type
		2	repeat	Repeat indicator
		3	mmsi	Maritime Mobile Service Identity
		4	status	Navigation status
		5	turn	Rate of turn
		6	speed	Speed over ground
6	data	7	accuracy	Position accuracy
		8	lon	Longitude
		9	lat	latitude
		10	course	Course over ground
		11	heading	True heading
		12	manoeuvre	Manoeuvre indicator
		13	spare	unused
		14	raim	RAIM flag
		15	radio	Radio status
7	padding	Number of fill bits to 6-bit boundary		
8	*xx	Checksum		

The GNSS and AIS packets are transmitted with UDP, following the NMEA 0183 interface standard which defines the industry standard sentence structure of messages (National Marine Electronics Association, 2024). Specifically, we consider RMC sentences from the GNSS, the recommended minimum specification for GPS data, with the structure given in Table 3.

The AIS receives AIVDM sentences of Position report class A, with the structure given in Table 4.

The ICS messages are transmitted over Modbus/TCP, communicated between the ICS and PLCs. We use two types of messages: *Read Holding Registers*, and *Write Single Register*. A read involves sending a request, roughly every second, to the PLC which responds with the current value held in the requested registers. In the case of the throttle, a read gives us the current engine RPM, while the rudder returns the current rudder angle. A write involves sending the request, for example, when the crew changes the throttle, and the PLC responds with confirmation. Therefore a write to the throttle contains the desired power level, and the rudder contains the target rudder angle. The format of these Modbus/TCP sentences is given in Table 5.

6.3. Visualisation architecture

Our visualisation is a web-based application, using the D3.js library to power the visual element (Bostock et al., 2011). We use a typical Model-View-Controller design pattern (Reenskaug, 1979). Data is fed into the application as JSONs of packet captures, which can be directly generated from PCAP files using Wireshark.[3] The input data is parsed, then processed in the data controller, filtering the data by time and grouping data as necessary for each component. The user can choose to speed up or slow down the speed at which the captures are played,

[3] https://www.wireshark.org/

Table 5
Modbus Structure.

Field	Structure	Description
1	sequence num	Sequence number
2	protocol id	Protocol ID (0 × 0000 = Modbus/TCP)
3	length	Packet length (bytes)
4	slave id	Slave address
5	function code	Modbus Function Code (0 × 03 = Read Holding Registers, 0 × 06 = Write Single Register)
6	data	[Read]: Address of (first) register requested (Request), or contents of requested registers (Response). [Write]: Address and value to write to register (Request), or address and value written (Response).

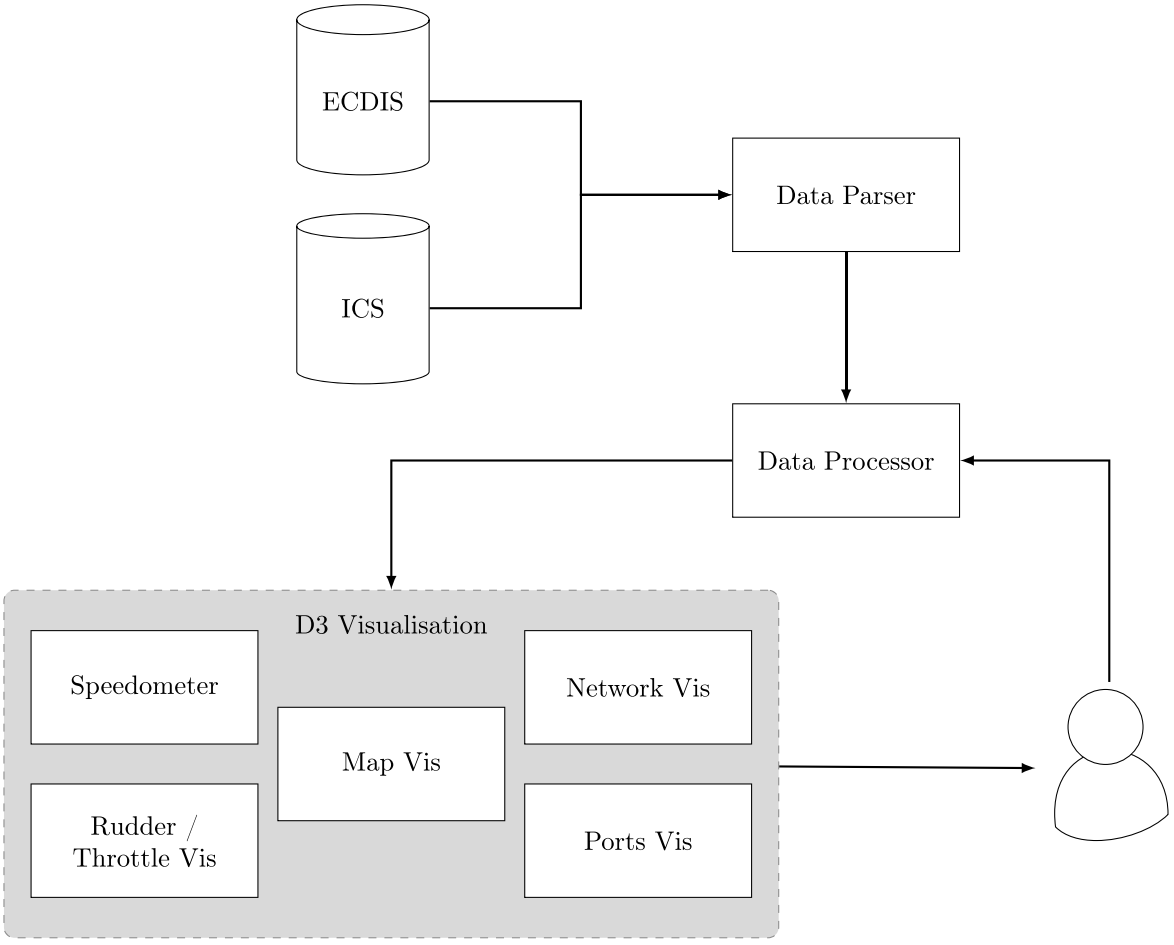

Fig. 5. Visualisation architecture.

can pause the playback, or skip to any point in time. A diagram of the architecture can be found in Fig. 5.

6.4. Graphical user interface

Our visualisation is composed of multiple coordinated views. It is designed with a stimulus-driven approach, with the intention of maximising visual saliency in the event of an attack. We also attempt to maintain aesthetics to attract the user and make the visualisation intuitive by taking cues from familiar visuals in maritime. In designing each visualisation, we minimise language-dependency, instead using iconography to convey information. This is to address one of the concerns raised in our interviews, regarding accessibility for users who may not necessarily speak English, and create a universally understood visual encoding. This also improves perception; Ware notes that cognition from an image generally occurs in a fraction of a second, much quicker than text (Ware, 2008). To minimise incorrect interpretations of icons, we use icons with high concreteness, according to Isherwood et al. to closely resemble the physical components that are represented (Isherwood et al., 2007), therefore improving the intuitiveness for the intended audience. We also utilise animations for changing elements. This allows the user to track changes, reducing cognitive load and mitigating change blindness. The user is also able to drag-and-drop each 'window' to enable efficient comparisons between views.

An overview of the visualisation can be seen in Fig. 6.

We anticipate that the Geospatial, Speedometer, and Throttle/Rudder views are familiar for crew members, whereas the network components are new concepts.

6.4.1. Geospatial

The primary view is the geospatial visualisation, plotting the ship's location received from the GNSS receivers, and the nearby ships received from the AIS transponder.

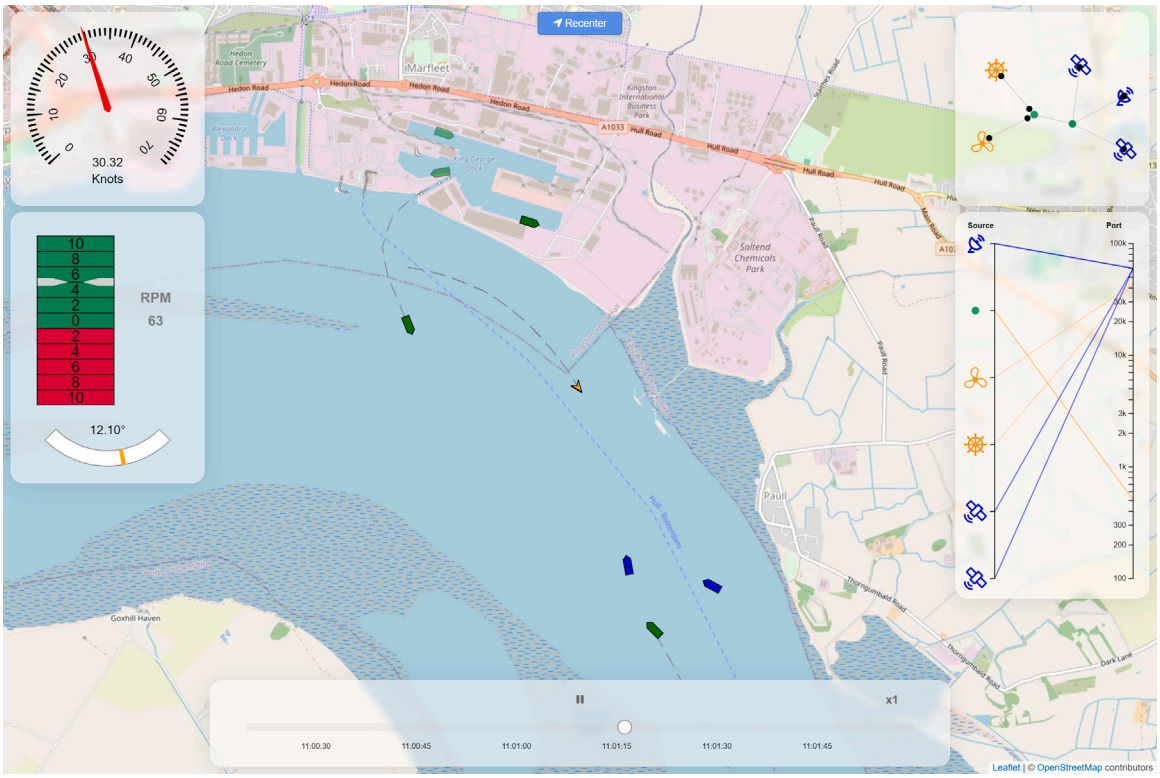

Fig. 6. Overview of the Visualisation Layout. The geospatial visualisation is designed to resemble traditional ECDIS charts.

Table 6
Geospatial Encoding.

Data Component	Visualisation Mapping	Description
Ship	Point	All signals from each ship are plotted on the map. Signals from the same ship (identified by MMSI) update the point.
Signal Source	Symbol	GPS → Arrow shape (orange); AIS → Ship shape. The GPS colour is chosen to be complementary to the sea, and distinct against other ships. The ship icon is designed to be reminiscent of the iconography found on a typical AIS system.
Ship Heading	Rotation	Heading → Symbol rotation.
Signal Newness	Colour Hue	New ship → Blue hue; Old ship → Green hue. Sequential colour scale indicates which ships we are seeing for the first time.
Signal Freshness	Opacity	New signal → Opaque; Old signal → Transparent. Opacity decreases linearly with time.
Longitude, Latitude	2D Position	Longitude, Latitude → Projected coordinates.
Track	Line	Each ship has a track that encodes its previous locations.
Distance Travelled	Length	Further travelled → Longer length.
Longitude, Latitude	2D Position	Longitude, Latitude → Projected coordinates of path.

From GNSS and AIS packets, we can extract the latitude, longitude, speed, course, among other fields by deep packet inspection, which we feed into our visualisation with the encoding in Table 6.

To avoid occlusion under a large number of ships, we ensure that our own ship identifier is always visible above others, and allow the user to highlight individual ships by hovering over them, with the ability to hide certain ships.

6.4.2. Speedometer

As previously, speed data is extracted from packets received from the GPS. The visualisation can be seen in Fig. 7, with the encoding described in Table 7.

For components analogous to cabin navigational equipment – the speedometer, throttle, and rudder visualisations – we take design inspiration from their physical counterparts, to invoke familiarity and improve intuitiveness.

6.4.3. Throttle/Rudder

The rudder and throttle data is extracted from packets sent between the ICS and the relevant PLC. Specifically, we obtain the current rudder angle from a read response, and the target rudder angle from a write request. Similarly, the engine RPM is obtained from a read response, and the power level is obtained from a write request. The visualisation can be seen in Fig. 8, with the encodings described in Tables 8 and 9.

Table 7
Speedometer Encoding.

Data Component	Visualisation Mapping	Description
Speed	Needle Angle	Speed → Angle on radial axis.
	Text Content	Speed → Text content. The speed is redundantly encoded by a text display, as typical on a ship's speedometer; this can be harder to quickly process for a human, but allows greater precision.
Change in Speed	Needle Motion	As the speed of the ship changes, the needle rotates to the appropriate position on the scale.

Table 8
Throttle Encoding.

Data Component	Visualisation Mapping	Description
Throttle Power	Needle Position	Power → Vertical position on power scale. The top half of the throttle, in green, indicates travel ahead, whilst the bottom half, in red, indicates astern. This colour scheme is typical on a ship.
RPM	Text Content	RPM → Text content.
Change in Power	Needle Motion	If the throttle is set to a different level, the needle moves to the appropriate position on the scale.

Table 9
Rudder Encoding.

Data Component	Visualisation Mapping	Description
Rudder Angle	Indicator Angle	Angle → Angle on radial axis.
	Text Content	Angle → Text Content. The angle is redundantly encoded by a text display for precision.
Target Angle	Colour Hue	Target angle → Grey hue; Rudder angle → Orange hue. The target angle indicates the angle that the crew have set the rudder to turn to.
Change in Angle	Indicator Motion	As the speed of the ship changes, the indicator rotates to the appropriate position on the scale.

Fig. 7. Speedometer Visualisation, designed to resemble traditional ship speedometers.

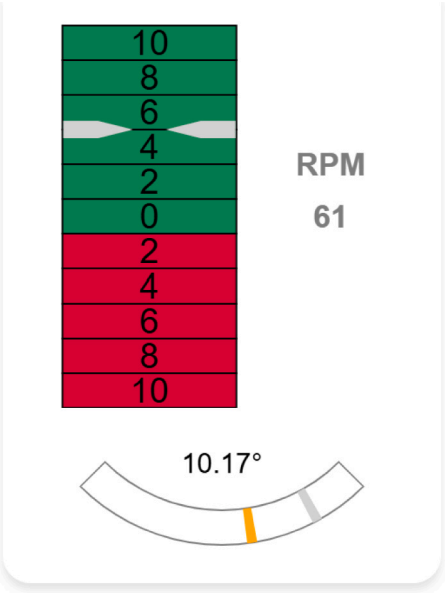

Fig. 8. Throttle and Rudder Visualisation, designed to reflect the appearance of their physical counterparts.

6.4.4. Network graph

The network graph provides an overview of the ship's network traffic. This data is extracted from the packet captures from both the ECDIS and ICS subnets, by inspecting packet headers to obtain the source and destination information. The visualisation can be seen in Fig. 9, with the encoding described in Table 10.

6.4.5. Ports graph

We use parallel coordinates to visualise the ports used in our ship's network, with two dimensions: the source IP, and the destination port. This information is retrieved from the packet headers of the network data. The visualisation can be seen in Fig. 10, with the encoding described in Table 11.

6.5. Attack scenarios

We have crafted 6 scenarios for the evaluation of our visualisation: 5 attacks and 1 'normal'. Each attack scenario is intended to show a feasible attack on a vessel, and demonstrate the use of the visualisation for detection. They are designed with inspiration from our literature review and attack taxonomy.

Table 10
Network Graph Encoding.

Data Component	Visualisation Mapping	Description
Endpoint	Node	Each endpoint is identified as a node in the graph.
Subnet	Colour Hue	ECDIS subnet → Blue hue; ICS subnet → Orange hue; Unknown → Red hue.
Endpoint Type	Icon	Specific components are given icons representative of their function, to be recognisable to the user.
Connection	Line	Connection between endpoints → Line between nodes.
Network Packet	Circle	Network packet → Circle on graph.
Packet Source	Start Position	Packet source → Circle position at respective node.
Packet Destination	End Position	Packet destination → Circle position at respective node.
Packet Travel	Motion	When the packet is observed, it moves from the source node to the destination node to indicate travel.

Table 11
Ports Graph Encoding.

Data Component	Visualisation Mapping	Description
Network Packet	Line	Network Packet → Line on graph.
Destination Subnet	Colour Hue	ECDIS Subnet → Blue hue; ICS Subnet → Orange hue. This categorical colour scheme remains consistent with the network graph.
Source IP	Dimension 1	Source IP → Vertical position on source axis. The iconography remains consistent with the network graph.
Destination Port	Dimension 2	Destination Port → Vertical position on port axis. This axis is overloaded for all destination IPs in our network, for simplicity in our visualisation to reduce need for interactivity to reveal information.
Packet Freshness	Opacity	New packet → 0.5 Opacity; Old packet → Transparent. Packets fade after a short period of time. A greater volume of traffic results in bolder lines.

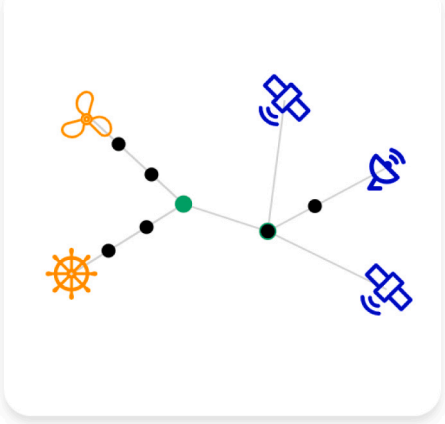

Fig. 9. Network Graph Visualisation, designed to convey the network architecture.

Scenario 0: Normal

The 'normal' scenario demonstrates a standard, non-attack scenario. We ensure that all components of the visualisation are included, such that this scenario is suitable to use for training purposes in our effectiveness study.

Scenario 1: Port scan

This first attack scenario shows a network reconnaissance attack: a port scan. Since it is a network attack, it is detectable on the ports graph and the network graph (Figs. 10, 9). This attack is modelled on a typical Nmap port scan, which probes the most popular 1000 ports (Lyon, 2008).

This could indicate a threat actor trying to map our network, or probing for weak entry points. The Industroyer is an example of malware targeting ICSs, which included a port scanner to map the target network (Cherepanov, 2017).

Scenario 2: Denial of service

This scenario deals with a denial of service enacted against the PLCs controlling the throttle and rudder. The ICS is infected with malware that is repeatedly sending control messages, impeding any legitimate commands. This is detectable on the network graph (Fig. 9), we can observe the control messages being received by the PLCs in the Throttle and Rudder visualisations (Fig. 8), and the effect on the ship's trajectory is observable on the map (Fig. 6).

This scenario is based on a similar attack theorised by Jones and Tam (2024). The consequences of such an attack can be inferred by that suffered after the 2021 accidental grounding of the Ever Given in the Suez Canal (Russon, 2021), as our scenario mimics this grounding as a consequence of the attack.

Scenario 3: GNSS spoofing

This scenario considers GNSS spoofing, where the normal GNSS signals are overpowered by signals on the same frequency containing false readings. This is detectable in our visualisation on the map, where we plot the GNSS signals (Fig. 6).

This particular spoofing scenario is a replication of behaviour observed in 2017 off the coast of Novorossiysk, where GPS receivers reported their location to be at an airport inland (Goward, 2017).

Scenario 4: GNSS jamming

Similar to the previous scenario, this attack scenario considers a GNSS jamming attack, where a threat actor emits radio signals to disrupt the low-power signals, preventing them from reaching the receiver. In our visualisation, we can observe the reduction in network traffic from the GNSS (Figs. 9, 10), the ships plotted on the map will fade as their most recent AIS signals age (Fig. 6), and our location will stop abruptly despite the speedometer (Fig. 7) and throttle (Fig. 8) indicating that we are in motion. We previously discussed the prevalent use of signal jammers in cargo thefts (Mukherji and Chandele, 2024); this scenario may be an indication of a pirate attack, with attempts to disrupt connections to the ship.

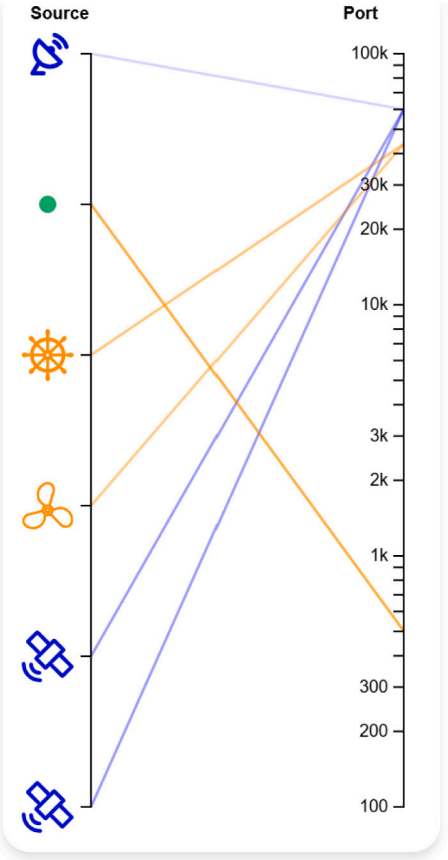

Fig. 10. Ports Graph Visualisation, designed to display network port activity.

Table 12
Participant Demographics: Cybersecurity Experience.

Cybersecurity Module	None
5	9

Scenario 5: Sybil

The final attack scenario is a Sybil attack, which has an attacker transmitting false AIS signals to create the illusion of traffic congestion. We can notice in our visualisation when the new signals are sent, by the increase in signals received by the AIS (Fig. 9), and the new plots on the map with their blue colour indicating newness (Fig. 6).

This scenario is inspired by the Sybil attacks that threaten autonomous vehicles (Rabieh et al., 2015). The impact of this kind of attack may be incorrect decision-making, particularly in the case of autonomous ships.

7. Effectiveness study

The study was carried out with 14 students, from a range of subject areas. The demographic of our participants is found in Table 12.

7.1. Quantitative analysis of study results

7.1.1. Detection accuracy

The proportion of true-positive attack detections are shown in Table 13.

Using the true-positive (tp), false-positive (fp), and false-negative (fn) detection rates, we calculate the precision, recall, and F-score as defined by Goutte and Gaussier (2005). The values for both tasks are presented in Table 14.

Table 13
Proportion of true-positive attack detections.

Attack Scenario	Pre-Training	Post-Training
1: Port Scan	3/3	14/14
2: DoS	3/3	14/14
3: GNSS Spoofing	3/3	14/14
4: GNSS Jamming	2/2	14/14
5: Sybil	3/3	14/14

Table 14
Precision, recall and F-score.

	Pre-Training			Post-Training		
	Precision	Recall	F-score	Precision	Recall	F-score
Cyber	0.83	1.00	0.91	0.89	1.00	0.94
Non-Cyber	0.82	1.00	0.90	0.87	1.00	0.93
Total	0.82	1.00	0.90	0.88	1.00	0.93

Table 15
Attack detection time [s].

Attack Scenario	Pre-Training		Post-Training	
	Mean	Standard Deviation	Mean	Standard Deviation
1: Port Scan	2.33	1.25	1.14	0.83
2: DoS	4.67	1.25	4.21	2.88
3: GNSS Spoofing	2.00	0.82	3.50	2.95
4: GNSS Jamming	8.50	2.50	6.93	3.61
5: Sybil	4.67	3.86	2.07	1.33
Overall	4.14	3.07	3.57	3.24

- Precision $p = \frac{tp}{tp+fp}$, the probability that a detection is correct given that the participant clicked.
- Recall $r = \frac{tp}{tp+fn}$, the probability that the participant clicked during an attack.
- F-score $F_1 = \frac{2pr}{p+r}$, the harmonic mean of precision and recall.

The post-training F-score (0.93) indicates that participants were able to accurately detect attacks, with the recall (1) showing that no attacks were missed, with a relatively low false-positive rate (0.88).

7.1.2. Detection efficiency

Fig. 12 shows the detection times of participants post-training, with red areas indicating the attack periods. Red bars indicate false-positive detections, and blue bars indicate true-positive detections. If participants clicked multiple times for the same reason, only the first is included.

From this graph, we observe that most participants were able to detect the attacks within a relatively short period. The mean and standard deviation for each attack scenario is given in Table 15, which shows that both pre- and post-training detection times were, on average, less than 5 s. GNSS jamming is a notable outlier here, due to its representation in the visualisation requiring participants to notice a lack of activity.

7.1.3. Intuitiveness

From Table 13, we see that in the pre-training task, all participants correctly detected the attack in each scenario. This might suggest that the visualisation is intuitive, and attacks are noticeable without training. However, based on responses when asked what happened in each scenario, participants were able to provide more detail and speculate on the consequences post-training, whereas pre-training they noticed anomalous behaviour but were often unable to explain it.

Figs. 11 and 12 show the detection times in pre- and post-training tasks respectively, with the mean detection times presented in Table 15. We observe that the mean detection time is not significantly different pre- and post-training.

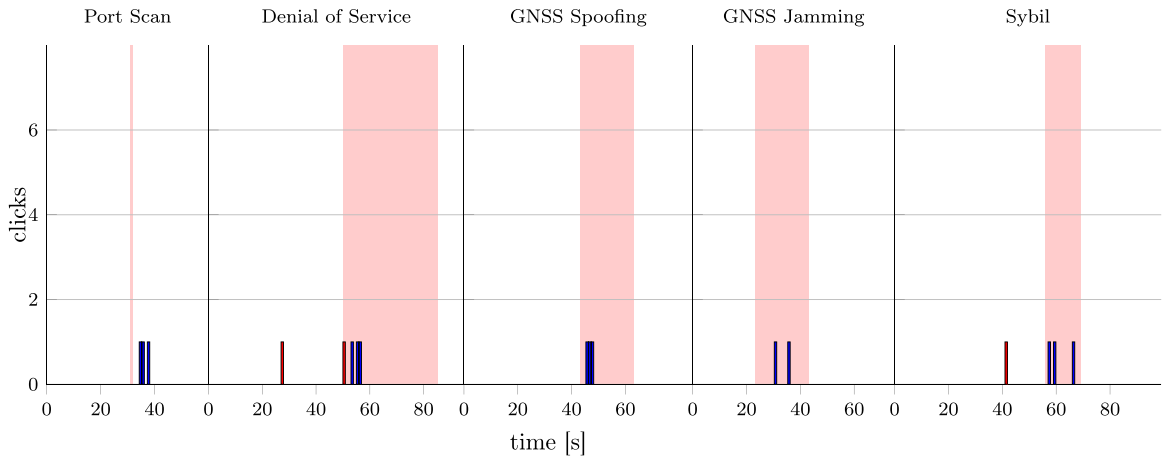

Fig. 11. Attack detection times in pre-training task.

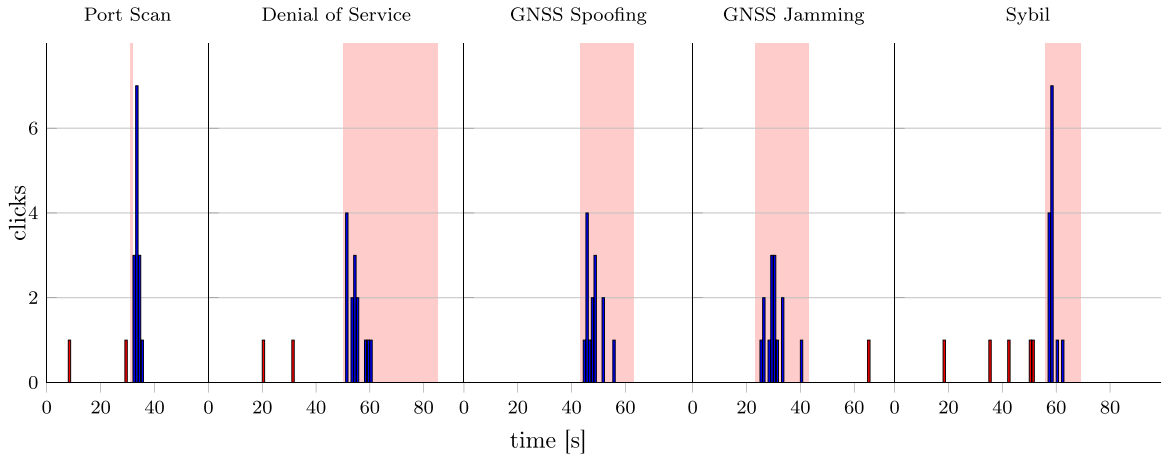

Fig. 12. Attack detection times in post-training task.

Table 16
Post-Training Attack detection time comparison.

	N	Mean μ	Std Dev σ
Cyber	25	3.32	3.47
Non-Cyber	45	3.71	3.09

However, the sample size of pre-training results is much smaller than that of post-training, so this result may not be significant. This is a limitation of our study method.

7.1.4. Effect of cybersecurity experience

In Table 14, we compare the precision, recall, and F-scores of participants who have and have not taken cybersecurity modules. The results show that the F-scores are similar regardless of whether they have studied cybersecurity.

We perform a two-tailed t-test for differences in mean detection times, with a 5% significance level.

Table 16 contains the relevant data. We hypothesise:

H_0 : Mean detection time is the same, $\mu_c = \mu_n$

H_a : Mean detection time is different, $\mu_c \neq \mu_n$

$$t = \frac{\mu_c - \mu_n}{\sqrt{\frac{\sigma_c^2}{N_c} + \frac{\sigma_n^2}{N_n}}} = -0.47$$

$$\mathbb{P}(|t| \geq 0.47) = 0.64 > 0.05$$

The results show that there was no significant difference in the mean detection efficiency between participants who have and have not studied cybersecurity.

7.2. Qualitative analysis of study feedback

We perform thematic analysis of the feedback. The codebook can be found in Appendix D.

Many participants commented positively on the aesthetics: *"I really like the interface, it's so nice and clean"*; *"It was very clear what each thing represents"*. One participant explained that because it is aesthetically pleasing, they felt compelled to *"try out different things and play around with it for longer periods of time"*.

Participants also remarked that the visualisation was *"easy to understand"*, noting that *"it's actually quite easy to notice something [abnormal]"*. One participant expressed that they had difficulty establishing a baseline for 'normal' behaviour of the physical ship components, *"how quickly it should turn, the average speed, engine power"*, highlighting a limitation of our study demographic not having maritime experience. However, they found that it was *"very easy to see some baseline for the network traffic"*.

Some participants reported that initially they felt overwhelmed by the amount of information presented: *"there's a lot of stuff I need to be looking out for"*, another commenting *"It was a little bit busy sometimes, so it felt like I might have missed something"*. However, after some time familiarising themselves with the visualisation, they found it easier to

Table 17
Interview Participant Demographics.

ID	Job Title	Sector
F1	Strategist	Security
G2	Consultant	Cybersecurity
G3	Consultant	Cybersecurity

understand: *"by using it more frequently, you get used to looking at all the things simultaneously"*.

The majority of participants commented that they were mainly drawn to the network components, with one participant stating that *"movement [made it] more compelling to look at"*. Another clarified that they used the network graph as a *"first indication of an attack . . . then I would check the engine, speed, and everything else"*.

All participants reflected that they felt more aware of threats after using the visualisation. Several participants noted that they were initially expecting more physical threats: *"I thought the ships were all going to start attacking each other"*; or not realising that there could be physical consequences: *"I thought it would just be communication attacks, I didn't think it could change the controls"*. Another participant commented *"all the examples were realistic, so they were quite enlightening"*.

Certain participants were more prone to false-positives. Largely, these false-positives were down to participants initially being unfamiliar with the visual encoding. Some inaccuracies can be ascribed to misunderstandings of the AIS, particularly the frequency of AIS messages compared to GPS messages, as well as difficulties establishing a baseline for operational technology. This can be attributed to the demographic's lack of maritime experience.

8. Evaluation

8.1. Qualitative analysis of evaluative interviews

We interviewed 3 participants with cybersecurity backgrounds and experience specific to maritime. 2 participants, G2 and G3, were previously interviewed in Section 4. Participant F1 leads the centre for cybersecurity at an international organisation, with previous experience as an officer onboard gas carriers. Participants G2 and G3 are both executives at a cybersecurity consultancy, which works closely with maritime organisations to provide emergency cyber-response, and cyber-training across the maritime ecosystem. G3 holds a position as an elected board member of an influential maritime organisation. The demographic of participants is found in Table 17.

Initial opinions of the visualisation were generally positive. F1 suggested that the visualisation gives a *"better understanding [of] interconnections between various pieces of infrastructure within a ship"*, with G3 explaining that it provides a good basis for demonstrating that system malfunctions might actually be targeted attacks. F1 further expanded that *"the consequences are really good, and the simulation in terms of what could happen"*.

The visualisation also has some intuitive elements. F1 explained *"it seems quite intuitive, I don't think it would be a challenge for [mariners] to understand"*. Specifically on the left-hand components: *"anybody who's in the shore-based world would understand that really, really easy"*. Similarly, G3 recognised the design elements from ECDIS and AIS. However, F1 raised some concerns over the network graphs requiring *"a little more of a narrative"* to make it *"idiot proof"*. G2 and G3 both shared doubts over the target audience. G2 explained that crew members may not care about the tool, as understanding the technical element is not their responsibility, with G3 emphasising that for an average, low-skilled sailer, this will not mean anything.

G3 suggested instead that the visualisation might be ideal from an insurance perspective, providing forensics to quickly determine the circumstances around incidents without necessarily requiring expert-level knowledge.

The visualisation has merit in training applications. G2 explained that bringing cyber into a training scenario is difficult, and the visualisation captures this well. G3 agreed, commenting that it adds to the bank of educational pieces. F1 suggested *"the appetite for training will probably be a bit higher"*, but as the tool is deployed, *"the possibilities might evolve"*. However, G2 commented that traction is difficult to attain, and uptake for solutions is generally very weak. G3 noted that if it were used only as a training tool, the scenarios wouldn't be replicable onboard ships without deploying similar network monitoring charts. Therefore, without being able to monitor the network activity in the way presented in the visualisation, crews would not be able to recognise the situations onboard from training with the tool. G3 also suggested that the tool might benefit from the inclusion of a playbook, describing how the crew should act in each scenario to mitigate the consequences. F1 similarly expressed that there was a *"step required in terms of what [the crew] could do"*.

8.2. Reflection

In response to this feedback, we might consider extending our visualisation to be more interactive for training applications; by allowing the user to perform actions, such as controlling the rudder and throttle, we can emphasise how the crew's inputs affect the attacks to create a more active learning environment. The network visualisation components, which are not usually available on a ship, could be displayed only after each scenario, as an explanation of what happened, thereby still allowing replicability onboard ships whilst building situational awareness of the underlying interconnectivity. We could then integrate a playbook, as suggested, to train the crew in appropriate incident response upon identifying a potential cyber event. These modifications would build a better narrative of the scenario, as well as increase interest in using the tool for the target audience.

9. Limitations

Our findings are subject to several limitations that must be acknowledged to contextualise their contributions.

A significant limitation lies in the participant demographics for the effectiveness study, which relied on a relatively small participant group of students without maritime experience. While this sample allowed us to capture initial insights into the tool's usability for non-experts, it is not reflective of the backgrounds or skill sets of maritime crew members. However, crew members onboard may have differing levels of technical and industry-specific knowledge which could influence their interactions with and perceptions of the tool, posing a potential challenge to the generalisability of our findings. The evaluative interviews aimed to reduce this limitation by capturing the maritime demographic. Whilst the participants have high levels of maritime experience, they are also highly trained in cybersecurity, and therefore neither the effectiveness study nor the expert evaluations fully capture the target demographic.

Additionally, the use of semi-structured interviews introduces inherent limitations. Whilst effective for gathering qualitative insights, interviews are limited by the subjectivity of participants and the potential for interviewer bias. By developing an interview protocol, we reduced the risk of interviewer bias, however the reliance on participants' self-reported experiences may not always accurately reflect broader trends within the maritime industry.

The use of simulation-based methods, while necessary due to the scarcity of real-world maritime cybersecurity data, also limits the applicability of our findings. While the synthetic data was designed to emulate realistic maritime scenarios, it cannot fully replicate the complexities, unpredictability, and variability of real-world maritime operations. This limitation is particularly significant as maritime environments often involve dynamic conditions and unforeseen challenges that are difficult to model. Consequently, the accuracy and efficacy

Table A.18
Notable Recent Maritime Cybersecurity Incidents.

Year	Description
2011	Port of Antwerp infected with malware by spear phishing. This enabled attackers to locate specific containers, find the security code for a container, change the location and scheduled delivery time, and smuggle out drugs before the scheduled pickup. After the infection was discovered, the same hackers broke into the facility and fit keystroke loggers to computers to continue their operations. (Direnzo et al., 2016)
2011	Iranian shipping line IRISL suffers a number of cyber-attacks, losing access to all data related to rates, loading, cargo number, date, and place. The attack also eliminated the company's internal communication network. This resulted in cargo being sent to the wrong destinations, causing severe financial losses. (CyberKeel, 2014)
2017	GPS interference reported by the U.S. Maritime Administration off the coast of Novorossiysk, Russia. It is suspected to be a GPS spoofing attack. (Goward, 2017)
2017	Danish shipping and logistics company Maersk falls victim to a malware attack, NotPetya. The attack disrupted global operations, costing the company $250–$300 million, and causing more than $10 billion in total damages. (Greenberg, 2018)
2018	Significant GPS interference reported in the Eastern Mediterranean Sea, resulting in lost or otherwise altered GPS signals affecting bridge navigation, GPS-based timing and communications equipment. (US Department of Transportation Maritime Administration, 2018a)
2019	US maritime facility taken down by Ryuk ransomware for over 30 h, after a phishing email is opened. Access to critical systems is lost. The same ransomware infects the Long Beach Port, and the ports of San Diego and Barcelona. (Cimpanu, 2019)
2019	Vessels in the Persian Gulf, Strait of Hormuz, and Gulf of Oman reported GPS interference, bridge-to-bridge communications spoofing, and/or other communications jamming. (US Department of Transportation Maritime Administration, 2018b)
2020	Port of Shahid Rajaee suffers a cyber-attack affecting 'computers that regulate the flow of vessels, trucks and goods', causing 'massive backups on waterways and roads'. (Beech, 2020)
2020	Port of Kennewick hit by ransomware, taking almost a week for port authorities to regain control over their data. The entry point is suspected to be a phishing email. (The Maritime Executive, 2020)
2022	Voyager Worldwide hit by a cyber-attack which took all systems down, affecting more than 20% of shipping companies worldwide. (Insurance Marine News, 2022)
2022	Port of Lisbon suffers a ransomware attack by LockBit, resulting in the capture of financial reports, company audits, budgets, contracts, cargo manifests, ship logs, information about crewmembers, personal data of customers, and port documentation, among other vital Port of Lisbon information. (The Maritime Executive, 2022)
2023	DNV suffers a ransomware attack affecting 1000 ships across 70 customers. It took two months for all users to be returned online. (Bergman, 2023)
2023	Port of Nagoya, Japan's biggest port, suffers ransomware attack by LockBit, forcing operations to be suspended for over a day. Toyota Motor Corporation subsequently had to suspend operations for more than 2 days following the attack. (Benjamin, 2023)
2023	DP World, which handles 40% of Australia's imports and exports, suffers a cyber-attack causing an enterprise shutdown for 3 days across Melbourne, Sydney, Brisbane, and Fremantle Ports. (Whitley and Doan, 2023)
2023	Significant GPS Interference and AIS Spoofing reported worldwide, with multiple instances occurring in the Strait of Hormuz. (US Department of Transportation Maritime Administration, 2023)

observed in controlled scenarios may differ in practical applications. Furthermore, the study evaluated only a small subset of attack examples. Real-world applications may involve complex, multi-vector attacks that were not considered in this study.

10. Conclusion

In this paper, we have designed and demonstrated the use of a cybersecurity visualisation for attack detection in a maritime context.

By conducting a literature review and interviews with participants across cybersecurity and maritime, we identified key cybersecurity challenges faced by the maritime industry, and design requirements for an effective visualisation to address these challenges. We proposed an attack taxonomy to further categorise attacks and inform the visualisation design. To demonstrate the use of the developed visualisation, we synthesised a dataset consisting of packet captures of a ship's network under realistic attack scenarios.

Unlike existing cybersecurity visualisations, which primarily focus on expert users, our approach is specifically tailored to the maritime domain, focussing on cyber-attacks against vessels. By designing the visualisation for non-expert users, we address a key industry gap where cyber situational awareness was often found lacking among mariners. Furthermore, the use of such non-expert visualisations could be employed to educate both non-technical operations and executive-level staff, to similarly raise cyber situational awareness at all levels and potentially improve the low levels of preparedness that many organisations suffer.

The results of our effectiveness study show that participants were able to use the visualisation we developed to accurately and efficiently detect synthetic attacks in an experimental setting. Qualitative results suggested an improvement in understanding both the threats and consequences of cyber-attacks on cyber–physical systems after using the visualisation under attack scenarios, and therefore an improvement in the cyber situational awareness of the users. The efficiency of attack detection was not significantly affected by cybersecurity experience, indicating that the visualisation is effective for non-expert users. This is particularly important given the capacity challenge and industry lack of cyber experts; we have developed a tool that can be used by a wider range of users to monitor and detect abnormalities, whilst also improving situational awareness. Results from evaluative interviews with maritime cybersecurity experts suggest that the visualisation has merit in being used as a training tool, and may have applications in digital forensics. Introduction of this tool in training environments could enable us to combat the industry-wide lack of awareness.

10.1. Future work

Following the promising results of our effectiveness study and feedback received from the expert evaluations, we might consider a longitudinal study to assess the visualisation over a longer period of time, in a production environment in the workflows suggested in the evaluation.

It will be necessary to scale the complexity of attacks. Our study and evaluation used relatively simple attacks, which yielded clear abnormalities in the visualisation. Going forward, we could assess the effectiveness against more subtle attacks, such as advanced persistent threats. We should also investigate scaling the complexity of the ship's network, and consider the impact this has on the effectiveness of the visualisation. This further refinement of the visualisation could also

Table C.19
Interview codebook.

Theme	Subtheme	Code	Description	Refs
Cybersecurity Challenges		Forensics	Difficulties surrounding the lack of data, or false data, driving actions and affecting post-incident investigations.	A1, B2, B3, C4
		IT vs OT	Differentiating between IT and OT systems.	A1
		Collaboration[a]	A lack of openness and collaboration across the industry, regarding any cyber-attacks or incidents.	B2, B3, D5
		Money[a]	Lack of funds allocated to cybersecurity issues.	B2, B3, C4
		Policy[a]	Problems relating to policies, including insurance.	A1, B2, B3, C4, E6
		Situational Awareness	A lack of situational awareness amongst staff and companies in general.	A1, B2, B3, C4
	Attacks	Malware[a]	Issues surrounding malware installation and impacting operations.	A1, B2, B3, C4, E6
		Phishing	Targetting humans through phishing emails.	B2, B3, D5, E6
		Ransomware	Challenges posed by malware disrupting operations, demanding a ransom.	B2, B3, C4
Training		Lack of Training	No, or minimal, cybersecurity training.	B2, B3, C4, D5, E6
		Basic Training[a]	A basic level of cybersecurity training (e.g. cyber-hygiene lecture).	A1, B2, B3, E6
		Cost[a]	Costs impacting the standard of cybersecurity training.	C4
		Effective Training	Training that is proved to be effective at raising the standard of cyber-awareness.	A1
Visualisation		Areas to Develop[a]	Areas about current visualisations that could be improved or developed.	A1
		Basic	Basic visualisations, such as KPIs or dashboards.	A1, B2, B3
		Complex[a]	Complex, effective visualisations.	A1
		None	No known instances of visualisations being used.	C4, D5
	Audience	Managerial[a]	Visualisations designed primarily for managerial or executive staff.	A1, B2, B3
		Crew[a]	Visualisations designed primarily for crew members onboard vessels.	A1
	Limitations	Accessibility[a]	Limitations caused by inaccessibility of visualisations.	A1, B2, B3
		Cost[a]	Cost implications, and lack of funding.	A1
		Different Audiences[a]	Limitations caused by different audience interests.	A1

[a] Induced code.

facilitate its potential to mitigate the challenge of physical isolation in maritime, by serving as an intermediary tool. This would enable non-technical crew members to investigate cyber-events without relying on shore-based operation centres, where connectivity cannot be guaranteed.

CRediT authorship contribution statement

Dominic Too: Writing – original draft. **Louise Axon:** Writing – review & editing, Supervision. **Ioannis Agrafiotis:** Writing – review & editing, Supervision. **Michael Goldsmith:** Supervision. **Sadie Creese:** Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Notable recent maritime cybersecurity incidents

See Table A.18.

Appendix B. Interview protocol

- Introduction — Share the purpose of the research, and aims of the interview.
- Ground rules — Collect informed consent, and address any questions or concerns.
- Questions —

1. Briefly describe your background.
   PROBE Maritime? Cybersecurity?
   PROBE Specific role?
2. From your experience, what are the key cybersecurity challenges faced?
   PROBE Key threats? Attacks?
   PROBE Novel to maritime?
   PROBE Challenges around data? Tooling? Experience?
3. How are these challenges overcome?
   PROBE Areas for improvement?
4. What is the current standard for cybersecurity training in maritime?
   PROBE Existent for non-technical staff?
   PROBE Where is it lacking?
5. Would you be able to describe, or share examples of visualisations that are commonly used, if any?

Table D.20
Study codebook.

Theme	Code	Description	Frequency
Aesthetics	Positive	Positive opinions of the aesthetics of the visualisation	6
	Negative	Negative opinions of the aesthetics of the visualisation	0
Usability	Easy to Use	Participant commented on the visualisation being easy to use	3
	Effective	Participant felt that the visualisation was effective for detecting attacks	14
	Clarity	Participant commented on how clear the visualisation was in conveying the state	7
	Overwhelming[a]	Participant found the visualisation overwhelming at some point	5
	Learning[a]	Participant became more comfortable using the visualisation after some use	3
	Establishing Normal[a]	Participant had difficulty establishing a normal baseline level for activity	1
	Operational Technology[a]	Participant had difficulty understanding the use of operational technology	1
Awareness	Improved Awareness	After the study, participant felt they had an improved situational awareness of maritime threats	14
	Physical Consequences[a]	Participant commented that they were previously unaware of the possibility of physical consequences	5
	Learning Signatures[a]	Participant commented that they felt they were learning the signatures of attacks as they use the visualisation	1
	Realistic[a]	Participant commented that the attack scenarios seemed realistic, making the attacks believable	3
Focus Point	General[a]	Participant did not have a specific focus point, but was monitoring each component	2
	Geospatial View	Participant mainly focussed on the geospatial view	4
	Network View	Participant mainly focussed on the network graphs	7
	OT View	Participant mainly focussed on the operational technology views	2
False-Positives	AIS[a]	False-positive detections because of misunderstandings of the AIS	2
	Operational Technology[a]	False-positive detections because of misunderstandings of the operational technology	2
	Other[a]	False-positive detections for other reasons that are not of note	6

[a] Induced code.

PROBE What kinds of tasks are these designed for? Monitoring, forensics etc.?

PROBE What are they telling us? Threats, or Vulnerabilities?

PROBE What data are they visualising? Or *not* visualising?

PROBE Are they suitable for their purpose? Limitations?

6. Under what contexts would different visualisations be useful?

PROBE What situations are visualisations used in?

PROBE What situations *could* visualisations be useful in?

PROBE Different users?

PROBE Difference between onshore vs offshore?

7. (If relevant) Do you have any cybersecurity data that you would be willing to share, that might be used to test a visualisation?

- Next steps — Share contact details for further communication, if necessary. Follow up on any agreed data sharing. Potential to share other interviewees' details for participation.

Appendix C. Interview codebook

See Table C.19.

Appendix D. Study codebook

See Table D.20.

Appendix E. Evaluation interview protocol

- Introduction — Share the purpose of the research, and aims of the interview.
- Ground rules — Collect informed consent, and address any questions or concerns.
- Demonstration — Explain and demonstrate the visualisation.
- Questions —

1. Briefly describe your background (if not previously interviewed).

PROBE Maritime? Cybersecurity?

2. What are your initial opinions of the visualisation?

PROBE Positive or Negative.

3. Do you think the visualisation would be intuitive for its intended demographic?

   PROBE Familiar design?

   PROBE Any problems?

4. Do you think the visualisation can be effective for detecting attacks?

   PROBE Does the design clearly communicate the state?

   PROBE Is abnormal behaviour recognisable, or does it require recall?

5. How do you think it would fit into current workflows?

   PROBE Training? Monitoring? Forensics?

Data availability

The data that has been used is confidential.

References

Accenture, 2015. Winning with the Industrial Internet of Things. Accenture, URL https://web.archive.org/web/20221101230041/https://www.accenture.com/t20160909T042713Z_w_/us-en/_acnmedia/Accenture/Conversion-Assets/DotCom/Documents/Global/PDF/Dualpub_11/Accenture-Industrial-Internet-of-Things-Positioning-Paper-Report-2015.pdfla=en. (archived: 2022-11-01).

Afenyo, Mawuli, Caesar, Livingstone D., 2023. Maritime cybersecurity threats: Gaps and directions for future research. Ocean & Coastal Management (ISSN: 0964-5691) 236, 106493. http://dx.doi.org/10.1016/J.OCECOAMAN.2023.106493.

Antonopoulos, Markos, Drainakis, Giorgos, Ouzounoglou, Eletherios, Papavassiliou, Giorgos, Amditis, Angelos, 2022. Design and proof of concept of a prediction engine for decision support during cyber range attack simulations in the maritime domain. In: 2022 IEEE International Conference on Cyber Security and Resilience. CSR, IEEE, pp. 305–310. http://dx.doi.org/10.1109/CSR54599.2022.9850280.

Aslan, Omer, Samet, Refik, 2020. A comprehensive review on malware detection approaches. IEEE Access (ISSN: 21693536) 8, 6249–6271. http://dx.doi.org/10.1109/ACCESS.2019.2963724.

Axon, Louise, Happa, Jassim, Goldsmith, Michael, Creese, Sadie, 2019. Hearing attacks in network data: an effectiveness study. Comput. Secur. (ISSN: 0167-4048) 83, 367–388. http://dx.doi.org/10.1016/J.COSE.2019.03.004.

Balduzzi, Marco, Wihoit, Kyle, Pasta, Alessandro, 2013. Hey captain, where is your ship? Attacking vessel tracking systems for fun and profit. Trend Micro, URL https://www.slideshare.net/trendmicro/captain-where-is-your-ship-compromising-vessel-tracking-systems. (Accessed: 14 November 2023).

Balduzzi, Marco, Wilhoit, Kyle, Pasta, Alessandro, 2014. A Security Evaluation of AIS. Trend Micro, URL https://documents.trendmicro.com/assets/white_papers/wp-a-security-evaluation-of-ais.pdf. (Accessed: 14 November 2024).

Banissi, Ebad, Forsell, Camilla, Marchese, Francis T., Johansson, Jimmy, 2014. Information Visualisation: Techniques, Usability and Evaluation. ISBN: 978-1-4438-5981-3.

Baraniuk, Chris, 2017. How hackers are targeting the shipping industry. BBC News, URL https://www.bbc.co.uk/news/technology-40685821. (Accessed: 14 November 2024).

Barriball, K. Louise, While, Alison, 1994. Collecting data using a semi-structured interview: A discussion paper. J. Adv. Nurs. (ISSN: 0309-2402) 19 (2), 328–335. http://dx.doi.org/10.1111/j.1365-2648.1994.tb01088.x.

Beech, Eric, 2020. Israel linked to cyberattack on Iranian port: Washington post. Reuters, URL https://www.reuters.com/article/us-mideast-iran-israel-cyber-idUSKBN22U363. (Accessed: 14 November 2024).

Benjamin, Jacob, 2023. OT cybersecurity breach disrupts operations at the Port of Nagoya, Japan. Dragos, URL https://www.dragos.com/blog/ot-cybersecurity-breach-disrupts-operations-at-the-port-of-nagoya-japan. (Accessed: 14 November 2024).

Bergman, Jamey, 2023. DNV: 'all users back online' two months after ShipManager cyber attack hit 1,000 vessels. Riviera, URL https://www.rivieramm.com/news-content-hub/news-content-hub/dnv-reports-cyber-attack-on-its-shipmanager-software-74466. (Accessed: 14 November 2024).

Blake, Tanya, 2017. Hackers took 'full control' of container ship's navigation systems for 10 hours. Resilient Navigation and Timing Foundation, URL https://rntfnd.org/2017/11/25/hackers-took-full-control-of-container-ships-navigation-systems-for-10-hours-ihs-fairplay. (Accessed: 14 November 2024).

Bolbot, Victor, Methlouthi, Oussama, Chaal, Meriam, Valdez, Osiris, BahooToroody, Ahmad, Tsetkova, Anastasia, Hellström, Magnus, Saarni, Jouni, Virtanen, Seppo, Owen, Douglas, et al., 2022. Identification and Analysis of Educational Needs for Naval Architects and Marine Engineers in Relation to the Foreseen Context of Maritime Autonomous Surface Ships, (MASS). Aalto University.

Bostock, Michael, Ogievetsky, Vadim, Heer, Jeffrey, 2011. D3 data-driven documents. IEEE Trans. Vis. Comput. Graphics (ISSN: 10772626) 17 (12), 2301–2309. http://dx.doi.org/10.1109/TVCG.2011.185.

Bothur, Dennis, Zheng, Guanglou, Valli, Craig, 2017. A critical analysis of security vulnerabilities and countermeasures in a smart ship system. In: The Proceedings of 15th Australian Information Security Management Conference. ISBN: 978-0-6481270-8-6, pp. 81–87. http://dx.doi.org/10.4225/75/5a84fe5695b55.

Braun, Virginia, Clarke, Victoria, 2006. Using thematic analysis in psychology. Qual. Res. Psychol. (ISSN: 14780887) 3 (2), 77–101. http://dx.doi.org/10.1191/1478088706QP063OA.

Byres, Eric, 2013. The air gap. Commun. ACM (ISSN: 00010782) 56 (8), 29–31. http://dx.doi.org/10.1145/2492007.2492018.

Caponi, Steven L., Belmont, Kate B., 2015. Maritime cybersecurity: A growing threat goes unanswered. Intellect. Prop. Technol. Law J. (ISSN: 15343618) 27 (1), 16–18.

Chang, C-H, Wenming, S, Wei, Z, Changki, P, Kontovas, CA, 2019. Evaluating cybersecurity risks in the maritime industry: A literature review. In: Proceedings of the International Association of Maritime Universities (IAMU) Conference. pp. 79–86, URL https://researchonline.ljmu.ac.uk/id/eprint/11929.

Cherepanov, Anton, 2017. WIN32/INDUSTROYER A new threat for industrial control systems. ESET, URL https://web-assets.esetstatic.com/wls/2017/06/Win32_Industroyer.pdf. (Accessed: 14 November 2024).

Cimpanu, Catalin, 2018. Ships infected with ransomware, USB malware, worms. ZDNET, URL https://www.zdnet.com/article/ships-infected-with-ransomware-usb-malware-worms. (Accessed: 14 November 2024).

Cimpanu, Catalin, 2019. US Coast Guard discloses Ryuk ransomware infection at maritime facility. ZDNET, URL https://www.zdnet.com/article/us-coast-guard-discloses-ryuk-ransomware-infection-at-maritime-facility. (Accessed: 14 November 2024).

CISA, 2015. Cobham sailor 900 VSAT buffer overflow vulnerability. Cybersecurity and Infrastructure Security Agency, URL https://www.cisa.gov/news-events/ics-alerts/ics-alert-15-030-01. (Accessed: 14 November 2024).

Corallo, Angelo, Lazoi, Mariangela, Lezzi, Marianna, Luperto, Angela, 2022. Cybersecurity awareness in the context of the Industrial Internet of Things: A systematic literature review. Comput. Ind. (ISSN: 0166-3615) 137, 103614. http://dx.doi.org/10.1016/J.COMPIND.2022.103614.

Creese, Sadie, Hannigan, Robert, El Kaafarani, Ali, Axon, Louise, Fletcher, Katherine, Schuler Scott, Arianna, Stolz, Marcel, 2020. Foresight review of cyber security for the industrial IoT. Lloyd's Register Foundation, URL https://www.lrfoundation.org.uk/publications/foresight-review-on-cyber-security-for-the-industrial-internet-of-things. (Accessed: 14 November 2024).

CyberKeel, 2014. Maritime cyber-risks. CyberKeel, URL https://web.archive.org/web/20150812150303/https://www.cyberkeel.com/images/pdf-files/Whitepaper.pdf. archived: 2015-08-12.

Demchak, Chris, Thomas, Michael, 2021. Can't sail away from cyber attacks: 'Sea-hacking' from land. War the Rocks, URL https://warontherocks.com/2021/10/cant-sail-away-from-cyber-attacks-sea-hacking-from-land. (Accessed: 14 November 2024).

Department of Homeland Security OIG, 2024. Coast guard should take additional steps to secure the marine transportation system against cyberattacks (OIG-24-37). URL https://www.oig.dhs.gov/reports/2024/coast-guard-should-take-additional-steps-secure-marine-transportation-system-against-cyberattacks/oig-24-37-jul24. (Accessed: 26 October 2024).

Direnzo, Joseph, Goward, Dana A., Roberts, Fred S., 2016. The little-known challenge of maritime cyber security. In: ISA 2015 - 6th International Conference on Information, Intelligence, Systems and Applications. Institute of Electrical and Electronics Engineers Inc., http://dx.doi.org/10.1109/IISA.2015.7388071.

DNV, 2023. Maritime cyber priority 2023. DNV, URL https://www.dnv.com/cybersecurity/cyber-insights/maritime-cyber-priority-2023. (Accessed: 14 November 2024).

Dyryavyy, Yevgen, 2014. Preparing for cyber battleships-electronic chart display and information system security. NCC Group, URL https://www.fox-it.com/media/cvtp5sqj/_2014-03-03_-_ncc_group_-_whitepaper_-_cyber_battle_ship_v1-0.pdf. (Accessed: 14 November 2024).

Endsley, Mica R., 1988. Design and evaluation for situation awareness enhancement. Proc. Hum. Factors Soc. Annu. Meet. 32 (2), 97–101. http://dx.doi.org/10.1177/154193128803200221.

Ericsson, 2023. Ericsson mobility report: June 2023. URL https://www.ericsson.com/49dd9d/assets/local/reports-papers/mobility-report/documents/2023/ericsson-mobility-report-june-2023.pdf. (Accessed: 14 November 2024).

Erstad, Erlend, Hopcraft, Rory, Vineetha Harish, Avanthika, Tam, Kimberly, 2023. A human-centred design approach for the development and conducting of maritime cyber resilience training. WMU J. Marit. Aff. 22 (2), 241–266. http://dx.doi.org/10.1007/s13437-023-00304-7.

EU, 2022. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the union. pp. 80–152, OJ L 333, URL https://eur-lex.europa.eu/eli/dir/2022/2555/oj. (Accessed: 14 November 2024).

EUSPA, 2021. Galileo open service navigation message authentication (OSNMA). European Union Agency for the Space Programme, URL https://www.euspa.europa.eu/newsroom-events/news/galileo-open-service-navigation-message-authentication-osnma-info-note-now. (Accessed: 14 November 2024).

EUSPA, 2023. ASGARD: The ultimate response to maritime spoofing attacks. European Union Agency for the Space Programme, URL https://www.euspa.europa.eu/newsroom-events/news/asgard-ultimate-response-maritime-spoofing-attacks. (Accessed: 14 November 2024).

Franke, Ulrik, Brynielsson, Joel, 2014. Cyber situational awareness - A systematic review of the literature. Comput. Secur. (ISSN: 0167-4048) 46, 18–31. http://dx.doi.org/10.1016/j.cose.2014.06.008.

Giles, Martin, 2019. Triton is the world's most murderous malware, and it's spreading. MIT Technol. Rev. URL https://www.technologyreview.com/2019/03/05/103328/cybersecurity-critical-infrastructure-triton-malware. (Accessed: 14 November 2024).

Goutte, Cyril, Gaussier, Eric, 2005. A probabilistic interpretation of precision, recall and f-score, with implication for evaluation. Lecture Notes in Comput. Sci. (ISSN: 03029743) 3408, 345–359. http://dx.doi.org/10.1007/978-3-540-31865-1_25.

Goward, Dana, 2017. GPS spoofing incident points to fragility of navigation satellites. Natl. Def. (ISSN: 00921491) 102 (766), 18–19, ISSN: 19433115, arXiv:27021938, URL https://www.jstor.org/stable/27021938. (Accessed: 14 November 2024).

Greenberg, Andy, 2018. The untold story of NotPetya, the most devastating cyberattack in history. Wired, URL https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world. (Accessed: 14 November 2024).

Hansman, Simon, Hunt, Ray, 2005. A taxonomy of network and computer attacks. Comput. Secur. (ISSN: 0167-4048) 24 (1), 31–43. http://dx.doi.org/10.1016/J.COSE.2004.06.011.

Harrell, Margaret C., Bradley, Melissa A., 2009. Data Collection Methods: Semi-Structured Interviews and Focus Groups. Technical Report, RAND Corporation, Santa Monica, CA.

Haugli-Sandvik, Marie, Lund, Mass Soldal, Bjørneseth, Frøy Birte, 2024. Maritime decision-makers and cyber security: Deck officers' perception of cyber risks towards IT and OT systems. Int. J. Inf. Secur. (ISSN: 1615-5270) 23 (3), 1721–1739. http://dx.doi.org/10.1007/s10207-023-00810-y.

Heering, D., Maennel, O.M., Venables, A.N., 2021. Shortcomings in cybersecurity education for seafarers. In: Developments in Maritime Technology and Engineering - Proceedings of the 5th International Conference on Maritime Technology and Engineering, MARTECH 2020, vol. 1, CRC Press/Balkema, pp. 49–61. http://dx.doi.org/10.1201/9781003216582.

Hopcraft, Rory, Harish, Avanthika Vineetha, Tam, Kimberly, Jones, Kevin, 2023. Raising the standard of maritime voyage data recorder security. J. Mar. Sci. Eng. 11 (2), 267. http://dx.doi.org/10.3390/jmse11020267.

Humphreys, Todd, 2013. UT austin researchers successfully spoof an $80 million yacht at sea. The University of Texas at Austin, URL https://news.utexas.edu/2013/07/29/ut-austin-researchers-successfully-spoof-an-80-million-yacht-at-sea. (Accessed: 14 November 2024).

IACS, 2023a. UR E26 REV1 CR: Cyber resilience of ships. International Association of Classification Societies, November 2023., URL https://iacs.org.uk/resolutions/unified-requirements/ur-e/ur-e26-new. (Accessed: 26 October 2024).

IACS, 2023b. UR E27 Rev1 CLN: Cyber resilience of on-board systems and equipment. International Association of Classification Societies, September 2023, URL https://iacs.org.uk/resolutions/unified-requirements/ur-e/ur-e27-rev1. (Accessed: 26 October 2024).

ICOM UK, 2024. What is AIS & how does it work? ICOM UK, URL https://icomuk.co.uk/What-is-AIS-and-How-Does-It-Work/3995/165. (Accessed: 14 November 2024).

ICS, 2021. The guidelines on cyber security onboard ships. International Chamber of Shipping, 4, URL https://www.ics-shipping.org/resource/guidelines-on-cyber-security-onboard-ships-version-four. (Accessed: 14 November 2024).

IEC, 2021. IEC 63154:2021; maritime navigation and radiocommunication equipment and systems - cybersecurity - general requirements, methods of testing and required test results. International Electrotechnical Commission.

Imperva, 2024. What is a Sybil attack? Imperva, URL https://www.imperva.com/learn/application-security/sybil-attack. (Accessed: 14 November 2024).

Inoue, S., Hirano, M., Kijima, K., Takashina, J., 1981. Practical calculation method of ship maneuvering motion. Int. Shipbuild. Prog. (ISSN: 0020868X) 28 (325), http://dx.doi.org/10.3233/isp-1981-2832502.

Insurance Marine News, 2022. Voyager worldwide reportedly hit by cyber attack. Insurance Marine News, URL https://insurancemarinenews.com/insurance-marine-news/voyager-worldwide-reportedly-hit-by-cyber-attack. (Accessed: 16 October 2024).

International Maritime Organization, 2004. Report of the maritime safety committee on its seventy-ninth session. International Maritime Organization, https://www.dco.uscg.mil/Portals/9/DCODocuments/MarineSafetyCenter/Tonnage/CommitteeDocs/MSC_79-23-Add.2_Report_of_the_MSC.pdf?ver=2017-06-20-121134-447. (Accessed: 14 November 2024).

International Maritime Organization, 2017. Resolution MSC.428(98) maritimg cyber risk management in safety management systems. IMO, https://wwwcdn.imo.org/localresources/en/OurWork/Security/Documents/ResolutionMSC.428(98).pdf. (Accessed: 14 November 2024).

Isherwood, Sarah J., McDougall, Siné J.P., Curry, Martin B., 2007. Icon identification in context: The changing role of icon characteristics with user experience. Hum. Factors: J. Hum. Factors Ergon. Soc. (ISSN: 00187208) 49 (3), 465–476. http://dx.doi.org/10.1518/001872007X200102.

Jacob, Stacy A., Furgerson, S. Paige, 2012. Writing interview protocols and conducting interviews: tips for students new to the field of qualitative research. Qual. Rep. (ISSN: 1052-0147) 17 (42), http://dx.doi.org/10.46743/2160-3715/2012.1718.

Jacob, Sarah, Roach, April, Baazil, Diederik, 2023. Pro-Russian hackers target website of Europe's largest port in rotterdam. Bloomberg, URL https://www.bloomberg.com/news/articles/2023-06-14/pro-russian-hackers-target-website-of-europe-s-largest-port-in-rotterdam. (Accessed: 14 November 2024).

Jacq, Olivier, Boudvin, Xavier, Brosset, David, Kermarrec, Yvon, Simonin, Jacques, 2019. Detecting and hunting cyberthreats in a maritime environment: specification and experimentation of a maritime cybersecurity operations centre. In: 2018 2nd Cyber Security in Networking Conference. CSNet 2018, Institute of Electrical and Electronics Engineers Inc., http://dx.doi.org/10.1109/CSNET.2018.8602669.

Jiang, Liuyue, Jayatilaka, Asangi, Nasim, Mehwish, Grobler, Marthie, Zahedi, Mansooreh, Babar, M. Ali, 2022. Systematic literature review on cyber situational awareness visualizations. IEEE Access 10, 57525–57554. http://dx.doi.org/10.1109/ACCESS.2022.3178195.

Jones, Steven, 2014. Addressing cyber security risks. Port Technol. Int. (62), 194–195, URL https://www.porttechnology.org/wp-content/uploads/2019/05/SAMI.pdf. (Accessed: 14 November 2024).

Jones, Kevin D., Tam, Kimberly, 2024. High impact malware targeting maritime infrastructure. In: Cavalcanti, Ana, Baxter, James (Eds.), The Practice of Formal Methods: Essays in Honour of Cliff Jones, Part I. Springer Nature Switzerland, Cham, ISBN: 978-3-031-66676-6, pp. 236–250. http://dx.doi.org/10.1007/978-3-031-66676-6_12.

Kalwar, Abhishek, Bhuyan, Monowar H., Bhattacharyya, Dhruba K., Kadobayashi, Youki, Elmroth, Erik, Kalita, Jugal K., 2020. Tvis: a lightweight traffic visualization system for DDoS detection. pp. 1–6. http://dx.doi.org/10.1109/ISAI-NLP48611.2019.9068666.

Kenney, Matthew, Macdonald, Fiona, 2023. Shifting tides, rising ransoms and critical decisions: progress on maritime cyber risk management maturity. CyberOwl, CyberOwl, URL https://cyberowl.io/wp-content/uploads/2023/10/CyberOwl_HFW_Thetius-Cyber-Security-Report-2023-Shifting-Tides-Rising-Ransoms.pdf. (Accessed: 14 November 2024).

Klahr, Rebecca, Shah, J.N., Sheriffs, P., Rossington, T., 2017. Cyber security breaches survey 2017. Department for culture Media and Sport, URL https://www.gov.uk/government/statistics/cyber-security-breaches-survey-2017. (Accessed: 14 November 2024).

Krishna, Ritika Raj, Priyadarshini, Aanchal, Jha, Amitkumar V., Appasani, Bhargav, Srinivasulu, Avireni, Bizon, Nicu, 2021. State-of-the-art review on IoT threats and attacks: taxonomy, challenges and solutions. Sustain. 2021, Vol. 13, Page 9463 (ISSN: 2071-1050) 13 (16), 9463. http://dx.doi.org/10.3390/SU13169463.

Latvala, Outi Marja, Keranen, Tommi, Noponen, Sami, Lehto, Niko, Sailio, Mirko, Valta, Mikko, Olli, Pia, 2017. Visualizing network events in a muggle friendly way. In: 2017 International Conference On Cyber Situational Awareness, Data Analytics And Assessment, Cyber SA 2017. Institute of Electrical and Electronics Engineers Inc., http://dx.doi.org/10.1109/CYBERSA.2017.8073400.

Lightner, Natalie, 2024. Buffer overflow attacks: Detection, prevention & mitigation. Synopsys, URL https://www.synopsys.com/blogs/software-security/detect-prevent-and-mitigate-buffer-overflow-attacks.html. (Accessed: 14 November 2024).

Lyon, Gordon, 2008. Nmap network scanning. Nmap.Org.

Maritime Security Alliance, 2024. Anti-boarding. Maritime Security Alliance, URL https://maritimesecurityalliance.com/anti-boarding. (Accessed: 14 November 2024).

McAfee, 2013. Visualizing a ddos cyber attack. McAfee, URL https://www.mcafee.com/blogs/internet-security/visualizing-a-ddos-cyber-attack. (Accessed: 14 November 2024).

McKenna, S., Staheli, D., Fulcher, C., Meyer, M., 2016. BubbleNet: A cyber security dashboard for visualizing patterns. Comput. Graph. Forum (ISSN: 1467-8659) 35 (3), 281–290. http://dx.doi.org/10.1111/CGF.12904.

Medina, Daniel, Lass, Christoph, Marcos, Emilio Perez, Ziebold, Ralf, Closas, Pau, Garcia, Jesus, 2019. On GNSS jamming threat from the maritime navigation perspective. In: FUSION 2019 - 22nd International Conference on Information Fusion. Institute of Electrical and Electronics Engineers Inc., http://dx.doi.org/10.23919/FUSION43075.2019.9011348.

Mekala, Sri Harsha, Baig, Zubair, Anwar, Adnan, Zeadally, Sherali, 2023. Cybersecurity for industrial IoT (IIoT): Threats, countermeasures, challenges and future directions. Comput. Commun. (ISSN: 0140-3664) 208, 294–320. http://dx.doi.org/10.1016/J.COMCOM.2023.06.020.

Meyer-Larsen, Nils, Müller, Rainer, 2018. Enhancing the cybersecurity of port community systems. Dyn. Logist. (ISSN: 21948925) 318–323. http://dx.doi.org/10.1007/978-3-319-74225-0_43.

Misas, J.P., Hopcraft, R., Tam, K., 2022. Future of maritime autonomy: cybersecurity, trust and mariner's situational awareness. Conf. Proc. ISCSS http://dx.doi.org/10.24868/10703.

Misas, J.D. Palbar, Hopcraft, Rory, Tam, Kimberly, Jones, K., 2024. Future of maritime autonomy: cybersecurity, trust and mariner's situational awareness. J. Mar. Eng. Technol. 23 (3), 224–235. http://dx.doi.org/10.1080/20464177.2024.2330176.

Mohanakrishnan, Ramya, 2022. Man-in-the-middle attack detection and prevention best practices. Spiceworks, URL https://www.spiceworks.com/it-security/data-security/articles/man-in-the-middle-attack. (Accessed: 14 November 2024).

Morgan Stanley, 2016. The Internet of Things and the new industrial revolution. Automation World, URL https://www.morganstanley.com/ideas/industrial-internet-of-things-and-automation-robotics. (Accessed: 14 November 2024).

Mukherji, Vivek, Chandele, A.K.S., 2024. GNSS jamming: An omnipresent threat. Geospatial World, URL https://www.geospatialworld.net/prime/special-features/gnss-jamming-an-omnipresent-threat. (Accessed: 14 November 2024).

Munzner, Tamara, Maguire, Éamonn, 2015. Visualization Analysis & Design. CRC Press, Boca Raton, Florida, ISBN: 0-429-08890-6.

National Marine Electronics Association, 2024. NMEA 0183. URL https://www.nmea.org/nmea-0183.html. (Accessed: 14 November 2024).

Nielsen, Jakob, Molich, Rolf, 1990. Heuristic evaluation of user interfaces. Conf. Hum. Factors Comput. Syst. - Proc. 249–256. http://dx.doi.org/10.1145/97243.97281.

Oruc, Aybars, Chowdhury, Nabin, Gkioulos, Vasileios, 2024. A modular cyber security training programme for the maritime domain. Int. J. Inf. Secur. (ISSN: 1615-5270) 23 (2), 1477–1512. http://dx.doi.org/10.1007/s10207-023-00799-4.

OWASP, 2021. A05:2021 security misconfiguration - OWASP top 10:2021. OWASP, URL https://owasp.org/Top10/A05_2021-Security_Misconfiguration. (Accessed: 14 November 2024).

OWASP, 2024. Buffer overflow. OWASP Found., URL https://owasp.org/www-community/vulnerabilities/Buffer_Overflow. (Accessed: 14 November 2024).

Patino, Luis, Ferryman, James, 2016. Semantic modelling for behaviour characterisation and threat detection. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Work. (ISSN: 21607516) 1282–1288. http://dx.doi.org/10.1109/CVPRW.2016.162.

Potamos, Georgios, Stavrou, Eliana, Stavrou, Stavros, 2024. Enhancing maritime cybersecurity through operational technology sensor data fusion: A comprehensive survey and analysis. Sensors 24 (11), 3458. http://dx.doi.org/10.3390/s24113458.

Potamos, Georgios, Theodoulou, Savvas, Stavrou, Eliana, Stavrou, Stavros, 2023. Building maritime cybersecurity capacity against ransomware attacks. Springer Proc. Complex. (ISSN: 22138692) 87–101. http://dx.doi.org/10.1007/978-981-19-6414-5_6.

Rabieh, Khaled, Mahmoud, Mohamed M.E.A., Guo, Terry N., Younis, Mohamed, 2015. Cross-layer scheme for detecting large-scale colluding sybil attack in VANETs. IEEE Int. Conf. Commun. (ISSN: 15503607) 7298–7303. http://dx.doi.org/10.1109/ICC.2015.7249492.

Reenskaug, Trygve Mikjel H., 1979. The Original MVC Reports. Technical Report, University of Oslo, URL http://urn.nb.no/URN:NBN:no-14314. (Accessed: 14 November 2024).

Russon, Mary-Ann, 2021. The cost of the Suez Canal blockage. BBC News, URL https://www.bbc.co.uk/news/business-56559073. (Accessed: 14 November 2024).

Santamarta, Ruben, 2015. Maritime security: Hacking into a voyage data recorder (VDR). IOActive, URL https://ioactive.com/maritime-security-hacking-into-a-voyage-data-recorder-vdr/. (Accessed: 14 November 2024).

Sasi, Tinshu, Lashkari, Arash Habibi, Lu, Rongxing, Xiong, Pulei, Iqbal, Shahrear, 2023. A comprehensive survey on IoT attacks: taxonomy, detection mechanisms and challenges. J. Inf. Intell. (ISSN: 2949-7159) http://dx.doi.org/10.1016/J.JIIXD.2023.12.001.

Scarfone, Karen, Mell, Peter, 2007. Guide to intrusion detection and prevention systems (IDPS). http://dx.doi.org/10.6028/NIST.SP.800-94, Special Publication (NIST SP), National Institute of Standards and Technology, Gaithersburg, MD.

Schinas, Orestis, Metzger, Daniel, 2023. Cyber-seaworthiness: A critical review of the literature. Mar. Policy (ISSN: 0308-597X) 151, 105592. http://dx.doi.org/10.1016/j.marpol.2023.105592.

Söner, Ömer, Kayisoglu, Gizem, Bolat, Pelin, Tam, Kimberly, 2023. Cybersecurity risk assessment of VDR. J. Navig. (ISSN: 0373-4633) 76 (1), 20–37. http://dx.doi.org/10.1017/S0373463322000595.

Spravil, Julian, Hemminghaus, Christian, von Rechenberg, Merlin, Padilla, Elmar, Bauer, Jan, 2023. Detecting maritime GPS spoofing attacks based on NMEA sentence integrity monitoring. J. Mar. Sci. Eng. 2023, Vol. 11, Page 928 (ISSN: 2077-1312) 11 (5), 928. http://dx.doi.org/10.3390/JMSE11050928.

Staheli, Diane, Yu, Tamara, Crouser, R Jordan, Damodaran, Suresh, Nam, Kevin, O'gwynn, David, Mckenna, Sean, Harrison, Lane, 2014. Visualization evaluation for cyber security: trends and future directions. In: Proceedings of the Eleventh Workshop on Visualization for Cyber Security. ISBN: 978-1-4503-2826-5, pp. 49–56. http://dx.doi.org/10.1145/2671491.

Storm, Darlene, 2013. Hack in the box: researchers attack ship tracking systems for fun and profit. Computerworld, URL https://www.computerworld.com/article/2475227/hack-in-the-box-researchers-attack-ship-tracking-systems-for-fun-and-profit.html. (Accessed: 14 November 2024).

Tam, Kimberly, Moara-Nkwe, Kemedi, Jones, Kevin D., 2021. The use of cyber ranges in the maritime context: Assessing maritime-cyber risks, raising awareness, and providing training. Marit. Technol. Res. 3 (1), 16–30. http://dx.doi.org/10.33175/mtr.2021.241410.

The Maritime Executive, 2020. Ransomware cripples IT systems of inland port in Washington State. The Maritime Executive, URL https://www.maritime-executive.com/article/ransomware-attack-cripples-systems-of-inland-port-in-washington-state. (Accessed: 14 November 2024).

The Maritime Executive, 2022. Cyberattack threatens release of port of Lisbon data. The Maritime Executive, URL https://maritime-executive.com/article/cyberattack-threatens-release-of-port-of-lisbon-data. (Accessed: 14 November 2024).

Tory, Melanie, Möller, Torsten, 2004. Human factors in visualization research. IEEE Trans. Vis. Comput. Graphics (ISSN: 10772626) 10 (1), 72–84. http://dx.doi.org/10.1109/TVCG.2004.1260759.

Tvergrov, Eli Anne, 2023. What do you do if a hacker takes over your ship? Partner Science Norway, URL https://partner.sciencenorway.no/ntnu/what-do-you-do-if-a-hacker-takes-over-your-ship/2174961. (Accessed: 14 November 2024).

UNCTAD, 2021. Review of maritime transport 2021. URL https://unctad.org/system/files/official-document/rmt2021_en_0.pdf. (Accessed: 14 November 2024).

United States Naval Academy, 2024. Ship Maneuverability, EN400: Principles of Ship Performance. Technical Report, United States Naval Academy, URL https://www.usna.edu/NAOE/_files/documents/Courses/EN400/02.09_Chapter_9-May20-.pdf. (Accessed: 14 March 2024).

US Department of Transportation Maritime Administration, 2018a. 2018-007-Eastern Mediterranean Sea-GPS interference. MSCI Advisory, URL https://www.maritime.dot.gov/msci/2018-007-eastern-mediterranean-sea-gps-interference. (Accessed: 27 October 2024).

US Department of Transportation Maritime Administration, 2018b. 2019-012-Persian Gulf, Strait of Hormuz, Gulf of Oman, Arabian Sea, Red Sea-Threats to Commercial Vessels by Iran and its Proxies. MSCI Advisory, URL https://www.maritime.dot.gov/msci/2019-012-persian-gulf-strait-hormuz-gulf-oman-arabian-sea-red-sea-threats-commercial-vessels. (Accessed: 27 October 2024).

US Department of Transportation Maritime Administration, 2023. 2023-005-Various-GPS interference & AIS spoofing. MSCI Advisory, URL https://www.maritime.dot.gov/msci/2023-005-various-gps-interference-ais-spoofing. (Accessed: 27 October 2024).

Škrlec, Zoran, Bićanić, Zlatimir, Tadić, Joško, 2014. Maritime cyber defense. In: 6th International Maritime Science Conference, Vol. 1. IMSC 2014, (ISSN: 1847-1498) pp. 19–25, URL https://trid.trb.org/view/1423294. (Accessed: 14 November 2024).

Ware, Colin, 2008. Visual objects, words and meaning. Vis. Think. 107–127. http://dx.doi.org/10.1016/B978-0-12-370896-0.00006-8.

Whitley, Angus, Doan, Lynn, 2023. Australia cyberattack leaves 30,000 containers stuck at ports. Bloomberg, URL https://www.bloomberg.com/news/articles/2023-11-12/australian-port-operations-slowly-resume-after-cyberattack-on-dp. (Accessed: 14 November 2024).

World Economic Forum, 2023. The global risks report 2023. World Econ. Forum (18), URL https://www.weforum.org/publications/global-risks-report-2023. (Accessed: 14 November 2024).

Wu, Mingtao, Moon, Young B., 2017. Taxonomy of cross-domain attacks on Cyber-Manufacturing system. Procedia Comput. Sci. (ISSN: 1877-0509) 114, 367–374. http://dx.doi.org/10.1016/J.PROCS.2017.09.050.

Yu, Xingjie, Guo, Huaqun, 2019. A survey on iIoT security. In: 2019 IEEE VTS Asia Pacific Wireless Communications Symposium. APWCS, Institute of Electrical and Electronics Engineers Inc., pp. 1–5. http://dx.doi.org/10.1109/VTS-APWCS.2019.8851679.

Zaghloul, M.S., 2014. Online ship control system using supervisory control and data acquisition (SCADA). Int. J. Comput. Sci. Appl. 3 (1), 6, URL https://web.archive.org/web/20170411115225/https://www.ij-csa.org//PaperInfo.aspx?ID=14328. archived: 2017-04-11.

Zelazny, K., 2014. Approximate method of calculating forces on rudder during ship sailing on a shipping route. TransNav, Int. J. Mar. Navig. Saf. Sea Transp. (ISSN: 2083-6473) 8 (3), http://dx.doi.org/10.12716/1001.08.03.18.

Zetter, Kim, 2014. An unprecedented look at stuxnet, the world's first digital weapon. Wired, URL https://www.wired.com/2014/11/countdown-to-zero-day-stuxnet. (Accessed: 14 November 2024).

Zhao, Hanning, Silverajan, Bilhanan, 2020. A dynamic visualization platform for operational maritime cybersecurity. In: Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). vol. 12341 LNCS, Springer Science and Business Media Deutschland GmbH, (ISSN: 16113349) pp. 202–208. http://dx.doi.org/10.1007/978-3-030-60816-3_23.

Dominic Too studied Computer Science at the University of Oxford, graduating with a Master's degree (MCompSci) in 2024. His research focussed on the cybersecurity of the maritime industry, including the development of a cybersecurity visualisation for intrusion-detection of cyber–physical systems by non-expert users.

Louise Axon is a Research Associate in Cybersecurity at the University of Oxford. Her research interests include network-security monitoring and intrusion-detection approaches, cyber risk and insurance, security and privacy of distributed ledger technologies, and cybersecurity capacity building. She holds a DPhil in Cybersecurity (Oxford), an MSc in Mathematics of Cryptography and Communications (Royal Holloway) and a BA degree in Mathematics and Music (Cardiff).

Ioannis Agrafiotis is a Research Fellow at the Department of Computer Science and James Martin Fellow at the Global Cyber Security Capacity Centre, University of Oxford. His research interests include capacity building in cybersecurity, risk analysis and resilience in the cyber domain, cyber insurance, and anomaly detection techniques. Ioannis holds a PhD in Engineering (Warwick), a MSc in Analysis, Design and Management of Information Systems (LSE) and a BSc in Applied Informatics (University of Macedonia, Greece).

Michael Goldsmith is Director of the Global Cyber Security Capacity Centre at the Oxford Martin School and a Senior Research Fellow in the Department of Computer Science and at Worcester College, University of Oxford. His research spans a wide range of topics within security, from the mathematical to the social. He received his DPhil in Computation from Oxford University three decades ago for work on support for specification logics, and has also worked in concurrency theory and formal verification through exhaustive state exploration.

Sadie Creese is Professor of Cybersecurity in the Department of Computer Science at the University of Oxford. She was founding director of the Global Cyber Security Capacity Centre at the Oxford Martin School and a member of the Coordinating Committee for the Cyber Security Oxford network. She is engaged in a broad portfolio of cybersecurity research spanning situational awareness, visual analytics, risk propagation and communication, threat modelling and detection, network defence, dependability and resilience and privacy.