

## Exploring Anomaly Detection for Marine Radar Systems

Antoine Saillard<sup>1,2</sup>[0000-0002-8376-2726], Konrad Wolsing<sup>1,2</sup>[0000-0002-7571-0555],  
Klaus Wehrle<sup>2</sup>[0000-0001-7252-4186], and Jan Bauer<sup>1</sup>[0000-0001-7631-6456]

<sup>1</sup> Cyber Analysis & Defense, Fraunhofer FKIE,  
Wachtberg, Germany {firstname.lastname}@fkie.fraunhofer.de  
<sup>2</sup> Communication and Distributed Systems, RWTH Aachen University,  
Aachen, Germany {lastname}@comsys.rwth-aachen.de

**Abstract.** Marine radar systems are a core technical instrument for collision avoidance in shipping and an indispensable decision-making aid for navigators on the ship’s bridge in limited visibility conditions at sea, in straits, and harbors. While electromagnetic attacks against radars can be carried out externally, primarily by military actors, research has recently shown that marine radar is also vulnerable to attacks from cyberspace. These can be carried out internally, less “loudly”, and with significantly less effort and know-how, thus posing a general threat to the shipping industry, the global maritime transport system, and world trade.

Based on cyberattacks discussed in the scientific community and a simulation environment for marine radar systems, we investigate in this work to which extent existing Intrusion Detection System (IDS) solutions can secure vessels’ radar systems, how effective their detection capability is, and where their limits lie. From this, we derive a research gap for radar-specific methods and present the first two approaches in that direction. Thus, we pave the way for necessary future developments of anomaly detection specific for marine navigation radars.

**Keywords:** Marine Radar Systems; Maritime Cyber Security · Intrusion Detection Systems · Anomaly Detection · Navico BR24

### 1 Introduction

Marine navigation has long benefited from marine radar technology, enabling navigators to accurately assert their position, heading, and distance to surrounding hazards. In recent decades, navigational instruments have followed global trends of digitization, with electronic chart display and information systems (ECDISs) superseding paper charts, communication equipment, Global Navigation Satellite System (GNSS) devices, and various sensors interfacing with each other over shipboard networks [8, 34], ultimately leading to increased autonomy of vessels [16, 18]. This trend likewise affects Marine Radar Systems (MRSs), where digital radar displays have replaced their analog predecessors and their networking with other navigation instruments allows them to display a unified view of the vessel’s state and surroundings to navigators [7].

Computer networks forming the digital backbone of modern vessels do not come without challenges. Alongside digitization, the maritime industry has seen increased interconnectivity, with readily available satellite and mobile communication networks allowing for remote maintenance, chart updates, and crew access to the Internet [34]. Thus, the previously “water”-gapped shipboard systems become increasingly connected to the outside world, leading to new cyber threats, a serious concern that the nascent field of maritime cybersecurity research seeks to address. In that regard, the vulnerabilities of fundamental maritime systems on board vessels have been intensively discussed [3, 5, 8, 12] and demonstrated for various maritime systems, such as the automatic identification system (AIS) [1, 4], GNSS [6], VSAT [26], and also radar [21, 33, 35], showing their possible impacts on the economy, ecology, and human lives. However, security research on MRSs remains limited, especially regarding techniques to effectively mitigate potential cyberattacks.

With regard to network-based cyberattacks, as for any communication network, securing MRSs’ networks consists of two complementary efforts: the *prevention* of attacks in the first place and the *detection* of successful intrusions. On modern vessels, the former is usually implemented through network segmentation, *i.e.*, the separation of crew and management IT systems from critical OT systems [14]. Unfortunately, the integrity of radar data exchanged also depends on this segmentation as the often proprietary communication protocols employed are largely unsecured, lacking authentication mechanisms [35].

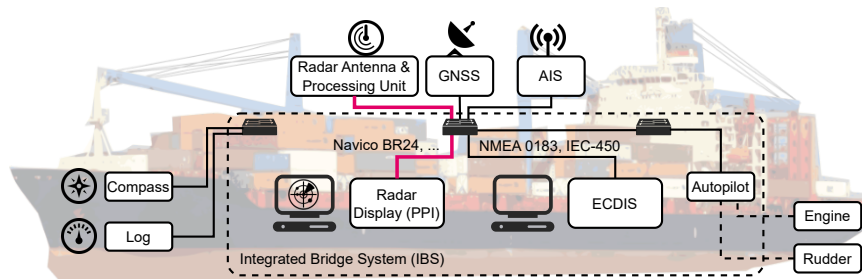
While preventive measures such as authenticated communication are desirable for the protection of MRS, they may be difficult to implement without the cooperation of radar vendors as modifications to the existing systems risks mandatory certifications to be voided. As complementary approach and topic of this publication, *detecting* manipulation of a radar image or changes in their communication in time for navigators to react is another viable option.

In that regard, the suitability of IDSs has long been researched for many domains [2, 32]. Yet, despite being a crucial component of shipboard navigation system networks [35], to date, detection mechanisms for MRSs have rarely been considered [22]. Thus, our work aims to address this research gap by investigating the applicability and evaluating the effectiveness of both existing and novel anomaly detection techniques for MRSs by making the following contributions:

- We update existing attack vectors on MRSs reported across different publications into a comprehensive and unified representation.
- To protect against these threats, we assess the effectiveness and capabilities of existing IDS solutions in a marine radar scenario.
- Based on identified limitations, we propose two novel image-based detection approaches to address the gaps in existing IDSs.

## 2 Background on Marine Radar Systems

On modern ships, an integrated bridge system (IBS) combines navigation equipment, controls over steering and propulsion, as well as communication devices



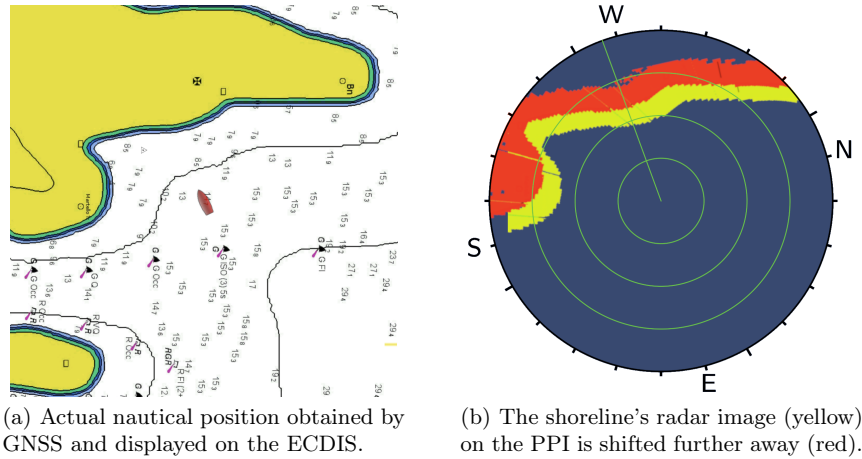
**Fig. 1.** MRSs, composed of a radar antenna unit and display, are connected via the IBS’s backbone network shared with other nautical equipment. Lacking proper preventive measures, the communication path, marked in red, is at risk.

in a unified system [3, 12] (*cf.* Fig. 1). Typical components include an ECDIS, radar, autopilot, GNSSs, speed logs, gyrocompasses, and AIS. To present a complete picture of the nautical situation to navigators, all these components are interconnected with the help of specialized maritime network protocols such as NMEA 0183 or its IP-based successors IEC 61162-450/460 [13, 14], which broadcast updates to the entire network, *e.g.*, regular GNSS position updates.

MRSs play a critical role in any IBS as they allow reliably determining the range and bearing to potential hazards, even at night or in adverse weather conditions [7]. Because of radar’s importance for collision avoidance, it is mandated by the International Convention for the Safety of Life at Sea (SOLAS) on passenger vessels and any vessel exceeding 300 gross tonnage [31].

Technically, MRSs can be abstracted into three main components [35], *cf.* Fig. 1: First, the transceiver antenna emits short directional electromagnetic pulses and receives the *echo*, *i.e.*, the signal reflection of objects. The measured time difference between the pulse’s emission and its echo’s reception allows for estimating the distance to the detected object. The echoes of all objects detected by orienting the pulse in a given direction through the rotating antenna constitute a *radar spoke*. Second, a processor unit performs the signal processing, typically also interconnected to the rest of the IBS, *e.g.*, to keep the image in a north-up orientation with data from the gyrocompass. Lastly, the radar display, also called Plot Position Indicator (PPI), is the interface for the operators that displays the radar image composed of the individual radar spokes and allows settings such as the scan range to be adjusted.

To transmit the radar spokes to the PPI, there exist many vendor-specific network protocols which share common features [35]. In this publication, we will focus on the Navico BR24 radar protocol [10] as it is a simple protocol that still exemplifies the characteristics of most radar communication protocols [35]. Its communication is split into three channels, each using a distinct UDP/IP multicast address: (i) the image channel carrying the raw radar image, (ii) the register



**Fig. 2.** Exemplary translation attack on the radar image shown by the PPI to the crew to let the shoreline appear further away than it actually is.

control channel carrying commands to the processor unit to adjust settings, and (iii) the report channel carrying meta-data from the processor unit [10].

Next, we concentrate on the most important communication channel, the image channel (i). BR24 subdivides the radar image into 2048 scanlines, each corresponding to a radar spoke with a span of about  $0.176^\circ$ . Scanlines are accumulated at the radar processor as generated by the turning radar and sent in aggregates of 32 scanlines per UDP packet. Each scanline carries a header with metadata such as the current angle, or the range of the scanline, as well as a payload of 1024 4-bit grayscale pixel values. Rendering each scanline pixel at the proper polar coordinates on the PPI yields the radar image over time. To adjust the image, parameters such as scan range, antenna rotation speed, and processing filters can be modified through the PPI's user interface, which sends the corresponding register messages to the processing unit. For more information, please refer to the publication by Dabrowski *et al.* [10].

Crucially, BR24 and other commonly used protocols lack essential security features, thus making maritime radar susceptible to cyberattacks [35]. Fig. 2 exemplarily depicts such an attack against the radar image. There, the image is shifted in the direction of travel so that the landmass on the PPI appears further away than it actually is. If carried out skillfully and situationally, *e.g.*, during poor visibility, this attack can compromise the safety of vessels.

### 3 Cybersecurity for Marine Radar

With the rising awareness of maritime cybersecurity, researchers have begun investigating the vulnerability of MRSs [21, 33, 35]. In our work, we consider

attacks that target the disruption of an MRS for the crew. An elaborated threat model is presented in the following Sec. 3.1. We then lay out different paths to remedy this current situation and state our research question in Sec. 3.2.

### 3.1 Threat Model

Executing attacks against MRSs is a non-trivial task. Referring to the MaCRA framework [34], we therefore assume any attacker to at least possess the resources of a *Tier*<sub>3</sub> threat actor. Such an attacker has the resources to exploit vulnerabilities to gain access to the IBS and the necessary domain knowledge and testing environments to develop specialized attacks, *i.e.*, understanding of protocols in use and the victim’s system as a whole. To more precisely specify the threat model MRSs are exposed to, we summarize related work by introducing five axes that distinguish the individual attack types.

**0) Attack Type.** As proven in related work, MRSs are susceptible to various attack vectors such as hardware- or software-based exploits [33], network attacks [35], and complementary threats from the domain of electronic warfare [21]. In 2020, Sviličić *et al.* conducted an extensive security assessment of radar systems on board two oil tankers [33], detecting many vulnerabilities in their underlying operating system, some of which were critical enough to enable a complete takeover of the system. Likewise, due to the typical lack of confidentiality and authentication mechanisms in the employed communication protocols, attackers may inject malicious commands or image data, as successfully demonstrated for Navico’s BR24 protocols [35] or ASTERIX CAT-240 [21, 22]. Such attacks can arbitrarily alter the displayed radar image on the PPI and deceive the crew’s understanding of the navigational situation (*cf.* Fig. 2).

**1) Attack Point.** Concerning network-based attacks, which are the focus of this publication, one overarching classification is the attacker’s position in the network, *cf.* red path in Fig. 1. In a Machine-on-the-Side (MotS) position, an attacker does not control the communication channel between processor and PPI, and can merely eavesdrop and inject packets. In a Machine-in-the-Middle (MitM) situation, the attacker has complete control over the communication channel and can prevent messages exchanged between processor and PPI from reaching their destination. The capabilities of MitM depend on the exact position of the attacker in the IBS topology. This publication considers an attacker isolating the PPI from the network with complete control over both the radar system’s communication and all the NMEA 0183 traffic exchanged with the PPI.

**2) Communication Channel.** MRSs have multiple communication channels (*cf.* Sec. 2), *i.e.*, the radar image stream, a control channel from the PPI to the radar unit, or a return channel for reports about the radar’s state. These channels can all be individually attacked. Generally, the implementation details of attacks against any channel heavily depend on the radar communication protocol used, the exact layout of the system, and the installed hardware.

**3) Manipulation Type.** Diving deeper into the attacks against the radar image, prior work distinguished between three main types of image manipulation [35]. *Denial of Service (DoS)* attacks deny using the radar image, *e.g.*, by

blinking the screen. *Transformation* attacks apply geometrical transformations to the image, affecting the bearing (rotation) or range (scaling) of plotted echoes. *Object manipulation* attacks alter the echoes of objects by (re)moving legitimate echoes or adding artificial ones, *e.g.*, to mimic buoys or other vessels.

**4) Stealthiness.** This axis concerns the degree to which the attack is conceived to remain undetected. On a technical level, mechanisms, such as source address spoofing, can be implemented to evade trivial protection measures. But, also the image manipulations itself can be conducted instantaneously, stealthily over time, or even be supported by external sensor information. For the latter, an attacker needs complementary information about the navigational situation to enhance the effect of their attacks [35]. This might include knowledge about visibility conditions, time of the day, or position of surrounding vessels reported by AIS to tailor an attack precisely to a challenging navigation situation. Another avenue is exploiting knowledge of the vessel’s speed and turn rate reported over NMEA 0183 to discreetly manipulate the radar image’s orientation during dynamic maneuvers, where minor discrepancies might be harder to notice [35].

### 3.2 Methods for Securing Marine Radar

Considering the spectrum of attacks, effective countermeasures are urgently needed. Optimally, preventive measures such as message authentication address the cause of the current threats. But in their absence or as additional security measure for defense-in-depth, we focus on detective measures in this publication.

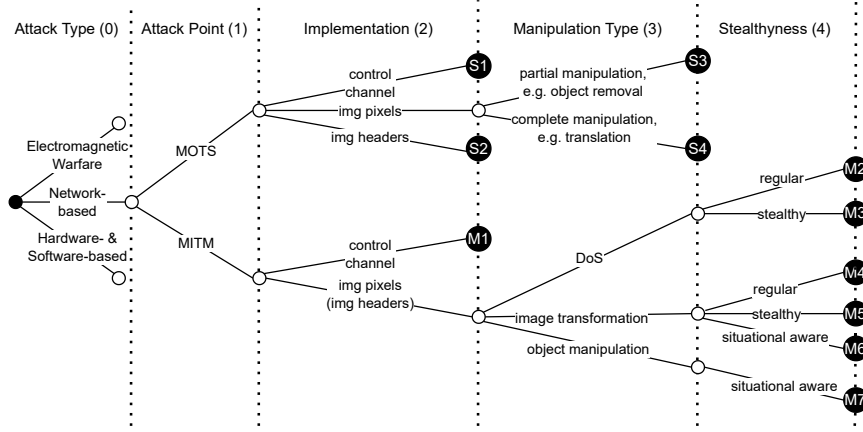
An IDS alerts the crew about (ongoing) attacks. Since past cyberattacks against MRSs are currently not documented or of an academic nature, we cannot assume that sufficient samples of attacks exist to train an IDS. Hence, we consider unsupervised IDSs trained on benign data to obtain a model of normality. Anomalies, *i.e.*, any deviation from these models, are considered an attack.

Naturally, there exists a vast body of commercial and academic IDS solutions. In this work, we consider four major classes: i) *rule-based* IDSs matching packets against a set of rules [30], ii) *timing-based* IDSs, *e.g.*, measuring packets inter-arrival times [20], iii) IDSs leveraging *machine-learning* to learn complex network patterns [23], and iv) IDSs that analyze the *application layer* data [22, 36], such as the radar image in our context. However, given the broad threat model, we believe that a single IDS is unlikely to adequately capture all attack vectors. Thus, our research question is:

*Which existing IDS solutions are best suited for the protection of MRSs and where need novel approaches to be developed in addition to fill in the gaps?*

## 4 Measurement Setup

To tackle the research question, we apply a selection of existing IDSs to marine radar network traffic and measure their capabilities. This section first details our considered attacks in Sec. 4.1. Then, we present the RadarSec-Lab environment and the IDS framework in Sec. 4.2 with which we conduct the experiments.




**Fig. 3.** Selection of network-based attack vectors, which we consider in this publication. We include four MotS attacks (S1–S4) and seven MitM attacks (M1–M7).

#### 4.1 Attack Vectors

Combining all axes from Sec. 3.1, a considerable number of attack vectors can be envisioned, especially because axis 4), stealthiness, is versatile. Thus, we restrict our analyses to a meaningful and reasonably sized subset, as listed in Fig. 3, and prune the attack tree accordingly. In the following, we briefly discuss our decisions for this selection.

For the attack point axis (1), we consider both variants, MotS and MitM. Regarding the implementation axis (2), our selection depends on the protocol that we analyze. Since Navico BR24 contains explicit control channels between the radar unit and PPI, *cf.* Sec. 2, we distinguish between hijacking the control channel and manipulating the image directly through the streamed pixels or indirectly through image headers. Note that we omit the latter when considering MitM attacks since manipulating a scanline’s angle in the packet header or on image-level results in the same effect on the PPI. The manipulation type axis (3) is only meaningful when considering image pixel manipulation. In the MotS case, we distinguish between attacks manipulating the whole image and those only affecting a small number of pixels, as they might influence communication timings differently. The complete distinction between all options is only made in the MitM case. Similarly, variants along the stealthiness axis (4) are only considered for MitM image manipulation to yield the most challenging scenarios. Here, we differentiate between regular stealthiness implementing just source address spoofing, stealthier attacks that apply modifications gradually over time, and even ones leveraging information exchanged via NMEA 0183 to be situational aware, *i.e.*, to rotate the image during a maneuver. These eleven attack classes cover a considerable extent of the possible attack vectors that promise to be challenging for IDSs to detect.



**Fig. 4.** We utilize the RadarSec-Lab [35] environment, the IPAL IDS framework [37], and Snort [30] to examine the effectiveness of IDSs in a marine radar environment.

#### 4.2 RadarSec-Lab Environment and IDS Framework

Given the attack definition, we still need tools to tackle our research question. First, a simulation environment to implement the derived attacks and second, a selection of IDSs for which we examine their effectiveness, *cf.* Fig. 4.

**Simulation.** While conducting our research with a real MRS system would be preferable, we decided for a simulation environment for scientific reproducibility and comparability in the future. The RadarSec-Lab was specifically designed to support marine radar cybersecurity research [35]. It simulates Navico BR24 radar communication by combining a customized version of the Bridge Command [25] ship simulator with OpenCPN [24], an open-source ECDIS. Finally, the Radar Attack Tool (RAT) [35] carries out various network-based attacks against BR24. In RAT, we added the capability also to conduct MotS attacks.

RadarSec-Lab generates a dataset for the training and evaluation of IDSs. We record one packet capture for each attack type, including a benign scenario without attacks (B). Descriptions of each attack can be found in Tab. 1 and at <https://zenodo.org/records/7188636/files/README.md>. Each capture lasts 1100 seconds and contains a randomized number of attacks with randomized duration with at least 70 seconds of benign traffic separating two consecutive attacks. Parameterizable attacks have randomized settings, *e.g.*, the angle by which the radar image is rotated or the time it takes to ramp up the rotation to increase stealthiness. Finally, to evaluate environment-specific effects, each scenario is run in three simulated worlds available in RadarSec-Lab (Simple Estuary, Santa Catalina, and Rostock) each with a different type of vessel and with three repetitions (six in the case of scenario B). The resulting dataset contains 117 packet captures, covering all 12 attack scenarios (B, S1–S4, and M1–M7) and totaling about 36 hours of simulation.

**Experiment.** The packet captures serve as input for the evaluation of various IDSs, including Snort [30] and IDSs of the IPAL IDS framework [37], which offers a plethora of validated (re-)implementation of established IDSs, *cf.* Fig. 4. To train them, we use one benign scenario of the Simple Estuary environment as the training dataset and the second benign repetition of Simple Estuary to adjust the IDSs’ sensitivity. Our configuration goal is to minimize the amount of false positives to avoid disrupting the crew during normal operations.

**Table 1.** Description of our attack scenarios provided by RAT.

Scen.	Description
B	MRS operates normally without attacks.
S1	Commands the radar to be turned off.
S2	Modifies the protocol header to rotate the radar image.
S3	Adds fake vessels and removes existing vessels from the radar image.
S4	Distort the radar image (random, blank screen, shift image).
M1	Commands the radar to be turned off.
M2	Distort the radar image (random, blank screen, static overlay)
M3	Freezes the radar image to the last received image.
M4	Scales, rotates, or shifts the radar image instantaneously.
M5	Scales, rotates, or shifts the radar image slowly over time.
M6	Scales, rotates, or shifts the image disguised in the vessels movements.
M7	Adds fake vessels and removes existing vessels from the radar image.

**Metrics.** To measure the IDSs’ effectiveness, we leverage three metrics. First, we consider an attack to be detected if an alarm is raised during its execution and state the fraction of successfully detected attacks. Second, the False Positive Rate (FPR) metric counts the number of false positives relative to the number of benign packets. Lastly, an IDS should also detect anomalous behavior, ideally in a timely manner. Therefore, we also consider the average time to first detection (TTD). Every metric is calculated individually for each of the twelve attack types but counting all three environments and three repetitions.

## 5 Effectiveness of Existing IDSs Adapted to MRS

We now assess which IDS direction of existing works proves promising in detecting which attack type. To this end, Sec. 5.1 introduces the four IDSs used in our analyses, and Sec. 5.2 shows their strengths and deficiencies in protecting MRSs.

### 5.1 IDS Selection

Research has proposed plenty IDSs to differentiate benign and malicious behavior. Given the four directions introduced in Sec. 3.2, we select one representative from each and present their core idea as well as how they transfer to MRS.

**Rule-Based (Snort).** A rule-based IDS, such as the prominent software Snort [30] considered here, monitors network traffic and inspects packets for patterns associated with known malware strains and attacks, matching their header and content against a set of rules. Unsurprisingly, a naive application of existing rules, *e.g.*, defaults included with the Debian Snort package, is ineffective in detecting any attack on our data. Thus, we exemplarily developed custom rules for Navico BR24 as depicted in Listing 1.1. The first two rules match packets with register commands disabling the radar unit. Rule 3 instead defines a constraint

```

alert udp any any -> any 6680 ( msg:"RadarOps A disabled";
  content:"|00 C1 00|"; depth:3; sid:6680;)
alert udp any any -> any 6680 ( msg:"RadarOps B disabled";
  content:"|01 C1 00|"; depth:3; sid:6681;)
alert udp any any -> any 6678 ( msg:"Duplicate scanline";
  content:"|00 44 0D 0E 00 00 34 92|";
  threshold:type both,track by_dst,count 2,seconds 1; sid:6678;)

```

Listing 1.1: Our Snort rules for Navico BR24. Rule 1 and 2 alert on setting the MRS to standby, while Rule 3 detects too frequent occurrences of a scanline.

on benign behavior. It assumes a maximum rotation speed on the radar antenna, expecting a scanline with a specific angle to be received not earlier than a certain time threshold (one second here) after the previous one.

**Timing-based (IaT).** A second IDS considers the inter-arrival time (IaT) between network packets, leveraging the periodicity of network traffic and raising an alarm when deviations exceed pre-determined thresholds [20]. To this end, IaT’s model estimates the mean  $\mu$  and standard deviation  $\sigma$  of the underlying IaT distribution. These values are the basis of the model’s upper and lower thresholds defined as  $\mu \pm N\sigma$ , where  $N$  is a user-defined sensitivity threshold, which we adjust to achieve no false positives on the training data. In the marine radar context, this IDS constructs a model for the radar image stream. Since the rotation speed of the radar antenna is constant, the network packets delivering the image step-by-step are assumed to occur regularly.

**Machine-learning (Kitsune).** From the domain of machine-learning, Kitsune is a widely used general-purpose network IDS [23]. It trains artificial neural networks to perform anomaly detection on network traffic. The IDS’s ability to apply to any UDP or TCP traffic makes it an interesting candidate. Internally, Kitsune keeps track of multiple features, such as the bandwidth or IaT, and calculates an anomaly score, the Root Mean Square Error (RMSE). The user defines a threshold upon which an alert is raised for each packet exceeding it.

**Application-layer (Steadytime).** The last IDS monitors application data, *i.e.*, the payload of network packets, which in our case contains the radar image and control commands. Steadytime [36] assumes that monitored features regularly change, such as the radar angle changing with each transmitted radar spoke. It measures the time a feature remains static and compares it to the minimal and maximal static time seen during training. More concretely, we select the following radar features: i) the scanline counter and angle values, as these should change regularly between each transmitted packet. ii) The sum of each image frame’s first scanline’s pixels and the sum of pixels in scanlines with angle 0, since given natural noise, perfectly static images can be considered anomalous.

## 5.2 Evaluation

Applying the IDSs to the recorded data, we obtain the results depicted in Tab. 2. First, we observe that all IDSs’ FPR is 0.00% if no attacks are present, *cf.* column B. This is interesting as we trained only on the Simple Estuary world but

**Table 2.** Performance of existing IDSs in a MRS context. No IDS raises a false alarm in the benign scenario (B), and they reliably detect the MotS attacks (S1–S4) yet struggle to detect MitM attacks, especially those targeting the radar image.

	Metric	B	S1	S2	S3	S4	M1	M2	M3	M4	M5	M6	M7
Snort	Det. Attacks [%]	–	100	100	100	100	0	2.2	2.3	4.7	2.2	0	7.1
	FPR [%]	0	0	0	0	0	0	0	0	0	0	0	0
	Mean TTD [s]	–	0	1.8	1.2	1.3	–	11.4	1.7	12.3	1.7	–	35.8
IaT	Det. Attacks [%]	–	0	100	100	100	0	0	0	0	0	0	0
	FPR [%]	0	9.0	4.5	4.3	4.3	9.1	0	0	0	0	0	0
	Mean TTD [s]	–	0	3.0	3.1	3.1	–	–	–	–	–	–	–
Kitsune	Det. Attacks [%]	–	100	100	100	100	100	8.9	2.3	0	2.2	0	0
	FPR [%]	0	29.3	6.6	6.2	6.2	27.4	0	0	0	0	0	0
	Mean TTD [s]	–	3.5	0.5	0.5	0.5	4.0	18.4	6.2	–	6.0	–	–
Steadytime	Det. Attacks [%]	–	0	100	100	100	0	60.0	100	7.0	2.2	9.1	4.8
	FPR [%]	0	0	0	0.1	0.1	0	1.5	1.6	0	0	0	0
	Mean TTD [s]	–	–	0	0.1	0	–	7.4	62.6	20.8	9.4	22.8	19.8


evaluated in all three worlds. We therefore conclude that no IDS is sensitive to environmental changes in terms of false positives. However, regarding the detected attacks, we find a clear distinction between MotS (S1–S4) and MitM (M1–M7).

The MotS attacks are detected remarkably well by all approaches, except for IaT and Steadytime failing to detect S1. Since S1 deactivates the radar and the IDSs emit the alerts too late namely when the radar is activated again. Next, the TTD is small with 0 to 3.1s, which is in the order of one full rotation of the radar (about 2.4s in RadarSec-Lab). During attack conditions, the FPR is generally higher than during benign behavior. While Snort and Steadytime still feature an FPR of 0 to 0.11% in MotS attacks, IaT and Kitsune reach levels of up to 29.3%. Both IDSs make use of window mechanisms, where past packets factor into the classification of the current packet, leading to trailing false positives.

This effect is particularly pronounced at the end of attacks in scenario S1, where the resumption of the image stream after a long interruption causes a notable jump in the IDSs’ anomaly scores. Fig. 5 depicts this phenomenon exemplarily for Kitsune. During attacks the anomaly score (blue curve) steadily rises in absence of packets. At the end of each attack when the radar image stream resumes, the jitter on that communication channel causes a notable jump in the anomaly score, leading to trailing false positives (yellow marker).

We deem these false positive rates to be acceptable due to the approaches’ excellent FPR results in the absence of attacks (B). Interestingly, the Snort rules outperform the more complex IDSs but require understanding the protocol’s packet format, and might fail in a more dynamic or noise-prone environment.

Despite these promising results, MitM attacks are, for the most part, not reliably detected. The purely network-based IDSs, Snort and IaT, fail to detect



**Fig. 5.** Performance of Kitsune on the S1 scenario. Attacks are highlighted in red, whereas alerts are indicated with purple (true alert) and yellow (false alert) markers. The blue curve depicts Kitsune’s anomaly score, and the black line is the alert threshold.

any attack sufficiently or just by chance. Kitsune can still detect M1, in which the traffic volume drops dramatically in the absence of the transmission of image data, as one of the features Kitsune monitors is the per-host bandwidth. Steady-time considering also the application layer, *i.e.*, the transmitted radar image, detects M2 and M3, where the radar image is static. But, this could trivially be circumvented by attackers by adding noise to the image.

**Takeaway.** Tying back to our research questions, existing IDSs are suitable for MRSs as they effectively detect MotS attacks while raising no false alarms during benign operations. However, it turns out that it is not feasible to detect the broad class of MitM attacks, especially those that change the radar image.


## 6 Radar-specific Approaches

Having identified a lack of detection capabilities for MitM attacks, we develop two novel IDSs designed to detect anomalies in the displayed radar image. First, we present and evaluate *Image-Delta*, a rather simple approach (Sec. 6.1), and then consider the more complex *Chart-Diff*, based on nautical charts (Sec. 6.2).

### 6.1 Radar Image-Delta IDS

According to the International Maritime Organization (IMO), MRS should scan their environment at a rate of no less than 12 rpm [15]. In comparison to vessel’s velocity, we assume that the difference between two consecutive full radar scans is small. Image manipulation attacks instead introduce noticeable and sudden changes, especially if they alter the entire image, *cf.* [35]. Hence, *Image-Delta* determines the amount of change, referred to as *delta*, between subsequent images and triggers an alarm if the observed change exceeds the acceptable range.

Because most of the difference between two consecutive images is expected to be caused by the vessel’s movement (translation and rotation), *Image-Delta* leverages NMEA 0183 reports for compensation, raising the sensitivity of the detector as shown in Fig. 6. The stabilization adds the current heading to each scanline’s angle to ensure a north-up orientation, *cf.* Fig. 6(b). Then, each pixel



**Fig. 6.** Stabilizing radar images with the vessel’s position and orientation from NMEA 0183 enables *Image-Delta* to remove legitimate differences in the images.

**Table 3.** *Image-Delta* outperforms existing IDSs regarding MitM but struggles in sophisticated attacks and performs worse in previously unseen environments.


Metric	B	S1	S2	S3	S4	M1	M2	M3	M4	M5	M6	M7
Det. Attacks [%]	–	0	44.4	0	47.7	0	75.6	61.4	100	31.1	15.9	7.1
FPR [%]	0	0.3	0.6	0	0.8	0.3	1.5	1.9	1.4	1.6	1.1	0.1
Mean TTD [s]	–	–	2.1	–	1.9	–	1.6	7.1	3.0	17.0	14.4	6.9

is offset by the recent movement of the vessel reducing even more legitimate differences between the figures, *cf.* fewer red echoes in Fig. 6(c). As a final step, we minimize the influence of noise with a Gaussian blur filter. The difference between all pixel values, *cf.* red parts in Fig. 6(c), is summed up and divided by the total value of pixels occurring in both images, giving a relative change between the images ( $\delta$ ). Scaling the delta in such a manner normalizes the final value, making it independent from radar specifics such as resolution, and reducing the influence of environmental factors such as the prevalence of landmasses.

Finally, this delta serves as the basis for the detection mechanism. During training on the Simple Estuary world, we measure the maximum  $\delta$ , which multiplied by a sensitivity parameter serves as the detection threshold. During detection, a  $\delta$  value exceeding the threshold causes an alert.

**Results.** We subject *Image-Delta* to the same evaluation as before, *cf.* Tab. 3. Again, we validate the absence of false positives across all benign scenarios *B*. The purely image-based IDS performs worse in MotS scenarios S1 and S3, where no or little manipulation of the image occurs. Still, substantial modifications to the image (S2 and S4) are detected. As *Image-Delta*’s goal is to address the shortcomings of existing IDSs, we focus on the MitM scenarios below.

Scenarios in which *Image-Delta* performs best are those where the image changes rapidly, *i.e.*, blanking or randomly overwriting the radar image (M2), freezing the image (M3), or suddenly scaling, rotating, or translating it (M4).



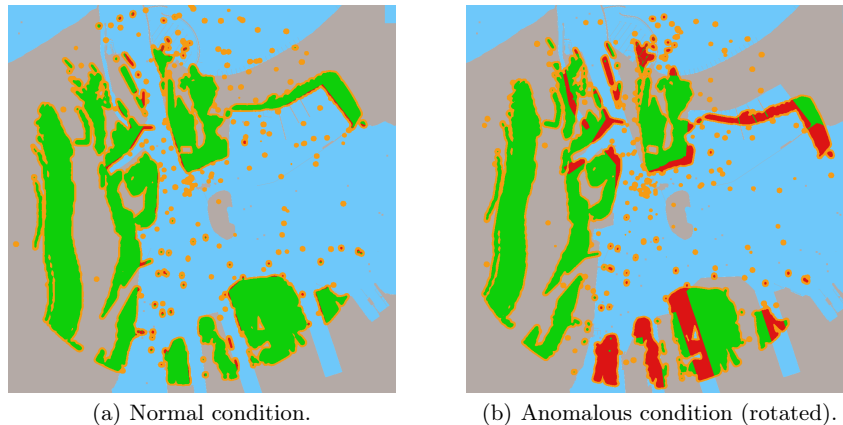
**Fig. 7.** Performance of *Image-Delta* on the M5 scenario. Attacks are highlighted in red, whereas alerts are indicated with purple (true alert) and yellow (false alert) markers. The blue curve depicts *Image-Delta*'s anomaly score, and the black line is the alert threshold. Note that the false positives at the end of attacks are caused by the sudden return to the unmodified image, resulting in a large delta in all cases.

The attack types less reliably detected are those that implement mechanisms to increase stealthiness, *i.e.*, applying image modifications slowly over time (M5) or scaled with the vessel's measured movements (M6). Note that switching off the radar (M1) is also not detected as no new radar images are generated. Lastly, missing the vast majority of attacks of M7, *Image-Delta* proves ineffective in this scenario. Only the echoes of single ships are manipulated, which is largely indistinguishable to the IDS from noise or inaccuracies in the stabilization and movement compensation process.

Concentrating on the stealthy attacks from M5 in the Simple Estuary world, which was used for training, the IDS detects 9 of the 15 attacks. But the detection performance is diminished on the other two scenarios (Santa Catalina and Rostock), *cf.* Fig. 7. *Image-Delta*'s performance depends on features of surrounding landmasses. Both benign changes and those caused by the attacker's gradual manipulation of the image are much more pronounced in Simple Estuary, causing the learned detection threshold to be inadequate for the two other environments. Depending on training conditions, the IDS can therefore be expected to underperform or cause false positives in unseen environments, a serious impediment to real-world applicability of the approach.

Lastly, we examine the real-time performance of *Image-Delta* as it has to handle large streams of uncompressed radar images. To this end, we measure the processing time for a single run of packet parsing and the detection methodology on a Intel i7-1365U CPU with 16GB of RAM. Parsing an 18.3 minutes long PCAP consisting of 362 724 Navico BR24 packets takes about 7.9 minutes, in large part due to the slow pyshark Python library. Our prototypical and unoptimized implementation of *Image-Delta* processes the parsed PCAP in 6.4 minutes. We therefore consider the approach to be feasible in processing Navico BR24 traffic in real time.

**Takeaway.** In contrast to previous IDSs, *Image-Delta* effectively detects sudden overlaying and transformation of the radar image even in MitM scenarios. It is significantly less effective in stealthier, gradual variants of the attacks, and inadequate in detecting object manipulation attacks targeting relatively small



**Fig. 8.** *Chart-Diff* counts the amount of invalid radar echoes (red) that do not correspond to the landmasses (gray). Weak echoes are attenuated through the filtering process and discarded as noise (orange), improving the sensitivity of the detector.

objects such as ship echoes. One drawback is that its performance can depend on terrain features, *e.g.*, the jagged landscape of Simple Estuary represents more opportunities for both benign inaccuracies and deliberate manipulation of the image to increase the anomaly score than the two other simulated world’s smoothly rising land. As a result of training *Image-Delta* on the Simple Estuary world, it is therefore not sufficiently sensitive for the two other environments.

## 6.2 Radar Chart-Diff IDS

Our second IDS aims to improve *Image-Delta*’s lack of environmental awareness. The main idea is to obtain some ground truth about the environment, *i.e.*, landmasses or buoys, from a trusted source. Since IBSs are obligated to have regularly updated, accurate digital charts on board for the areas in which the ship operates, these can be correlated with the radar image. Hence, the idea is to overlap the radar image with the charts and measure their difference.

For that purpose, *Chart-Diff* considers each radar echo independently. If an echo’s position corresponds to a location that is marked as a landmass on the chart, this pixel is considered valid, *cf.* green pixels in Fig. 8(a). Note that not every pixel on the radar image representing landmasses must contain an echo since landmasses may be shadowed, *e.g.*, through mountains. In contrast, invalid (non-verifiable) echoes can fall into four categories for which ground truth is hard to obtain: other vessels, uncharted objects, inaccuracies, or noise. Although information about other vessels can be obtained from AIS, not all vessels are equipped with AIS. Also, AIS signals could be tampered by sophisticated attackers [4], and indeed RAT already manipulates them appropriately [35]. Likewise, uncharted

**Table 4.** *Chart-Diff* is capable of detecting sophisticated manipulations such as translation, or rotation, by having more information available about the area from charts.

Metric	B	S1	S2	S3	S4	M1	M2	M3	M4	M5	M6	M7
Det. Attacks [%]	–	0	100	0	61.4	0	46.7	95.5	81.4	84.4	56.8	2.38
FPR [%]	0	0	0	0	0	0	0	0	0	0	0	0
Mean TTD [s]	–	–	4.6	–	3.9	–	3.2	25.4	7.3	26.2	20.8	9.7

objects or remaining inaccuracies cannot be predicted. Still, we apply a Gaussian blur filter to the image to find and ignore small or weak echoes resulting from noise, *cf.* orange pixels in Fig. 8(b). For detection, *Chart-Diff* learns the maximum proportion of invalid pixels, *i.e.*, red pixels, during normal conditions. If an attack tampers with the radar image, *e.g.*, rotating it, we expect the number of invalid pixels to increase and exceed the normal maximum, *cf.* Fig. 8(b).

**Results.** Like *Image-Delta*, *Chart-Diff* is also unable to detect S1, S3, and M1, but performs better in S4, M3, M5, and M6, especially in the scenarios we aimed to improve, such as M5, where the detection rate increases from 31.1% to 84.4%. *Chart-Diff* even makes great progress on M6, containing the stealthiest image manipulations conducted during vessel maneuvers, *e.g.*, rotating the radar image while the vessel is turning. Lastly, leveraging ground truth from the charts reduces environment-specific influences on the anomaly score: unlike *Image-Delta*’s, surrounding landmasses’ feature may vary widely as long as the accuracy of the charts remains comparable. In scenario M6 we nonetheless find *Chart-Diff* to underperform in the Santa Catalina world. There, the vessel’s few course changes and the relatively short duration of attacks do not provide enough opportunity for the attacker to cause significant rotation of the image. Translation in the direction of travel is also largely undetected as the vessel moves parallel to a relatively straight shoreline of this world.

We again measure the computational performance of our newly proposed IDS with the same methodology as for *Image-Delta*, *cf.* Sec. 6.1. *Chart-Diff* takes 5.7 minutes to process the 18.3 minutes long PCAP thereby being slightly faster than *Image-Delta*. Again, this IDS proves promising for time detection in real-time.

**Takeaway.** As an alternative radar-specific IDS, *Chart-Diff* outperforms *Image-Delta* on stealthier attacks. Additionally, one advantage of *Chart-Diff* is that its alerts can be visualized to the crew, *e.g.*, by highlighting suspicious echos once an alert is raised, *cf.* Fig. 8. Thereby, navigators can validate and reflect the ongoing situation and assess it themselves if necessary.

## 7 Discussion

Closing our analyses, we present related work and summarize the state of intrusion detection for MRS in Sec. 7.1. Finally, we discuss limitations in Sec. 7.2.

**Table 5.** Our evaluation on the effectiveness of IDSs reveals that MotS attacks are reliably detectable by established IDSs while lagging in MitM scenarios. Set out to remedy this situation, *Image-Delta* and *Chart-Diff* make progress in that direction, yet still leave room for improvement in the most sophisticated attacks (M6 and M7).

IDS	S1	S2	S3	S4	M1	M2	M3	M4	M5	M6	M7
Snort [30]	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
IaT [20]	✗	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
Kitsune [23]	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
Steadytime [36]	✗	✓	✓	✓	✗	(✓)	✓	✗	✗	✗	✗
Longo <i>et al.</i> [22]	✗	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
<i>Image-Delta</i>	✗	(✓)	✗	(✓)	✗	✓	✓	✓	(✓)	(✓)	✗
<i>Chart-Diff</i>	✗	✓	✗	(✓)	✗	(✓)	✓	✓	✓	(✓)	✗

The results by Longo *et al.* [22] are derived from their publication and not measured by us.

## 7.1 Related Work and Lessons Learned

Detecting anomalies in radar data caused by cyberattacks has been recently tackled in the aerospace sector [9, 19, 28, 29]. Yet, these works conceptually differ from challenges of the maritime domain as they consider (pre-)processed flight paths of tracked objects in a stationary setting, whereas vessels are constantly changing their location. Beyond tracking ships, navigators are likewise interested in perceiving landmasses. In the maritime domain, radar images combined with charts have been utilized for localization and detection of GNSS tampering [11], similar to *Chart-Diff*, but with a different purpose. The only directly related work we found is presented by Longo *et al.* [22], who propose a policy-based IDS for the ASTERIX CAT-240 radar protocol in a maritime scenario. Their policies enforce a constant rotation speed and a strict monotonicity of message identifiers, which is analogous to detecting the maliciously injected duplicates of image frames in the MotS scenarios S2–S4.

In a final assessment, shown in Tab. 5, we summarize the effectiveness of each IDS for the individual attack vectors. In addition to our experiments, we take the results of Longo *et al.* [22] into account, yet marked as grey since we transferred their results to our attack taxonomy without extra evaluations. Concerning our analyses of existing approaches and related work (upper and middle part of Tab. 5), we conclude that MRS can be effectively protected against MotS attacks using existing IDS approaches when adapted to radar network traffic. Set out to address the gap in MitM attacks, novel radar-specific, image-based IDSs such as *Image-Delta* or *Chart-Diff* provide a remedy in these situations especially for attacks M2–M5. But they still struggle with stealthy attacks that are carefully executed by sophisticated adversaries. Attacks modifying individual echoes, such as a nearby vessel (M7), are currently not covered by any approach. In conclusion, considering the pros and cons of each approach considered in our evaluation, currently, only a combination of multiple IDSs promises effective protection MRSs against most attacks.

## 7.2 Limitations and Future Work

As a simulation environment, RadarSec-Lab offers the possibility of scientifically reproducible experiments, but falls short in replicating user-induced behavior of real seafarers. Thus, we have not yet considered such behavior, *e.g.*, changing the radar resolution or scanned ranges. Our environment also has two other technical limitations. First, while the network traffic patterns are largely comparable to a real network, RadarSec-Lab supports only one of many network protocols (Navico BR24). Still, according to prior work, MRS protocols from different vendors are relatively similar [35]. Second, the radar echoes simulated by BridgeCommand lack realism, since compared to real MRS they are of lower resolution, overly sharp, and only approximate physical effects such as reflectivity.

With regard to a deployment of Radar-IDSs, future work still faces challenges. First, we trained and evaluated with little data compared to vessels that operate for weeks. Second, beyond detection performance, giving appropriate advice to the crew about the MRS's state is crucial to be of actual benefit, as proposed in [27]. Here, *Chart-Diff* IDS may be a first step in that direction, *cf.* Fig. 8. Next, the issue of verifying other vessels and AIS signals persists. While Katsilieris *et al.* [17] have validated AIS signals with radar echoes, nowadays both sources have to be considered vulnerable. Finally, the most significant conceptual issue of image-based IDSs is that these are mainly suited for coastal regions. In open waters, where there are virtually no reference signals to determine whether the image is altered, detecting anomalies is much more difficult. This scenario as well as coordinated attacks simultaneously targeting multiple IBS components, *i.e.*, radar, AIS, and NMEA 0183, are a challenging field for future research.

## 8 Conclusion

Modern vessels' integrated bridge systems (IBSs) and thus also their radar systems have experienced increased interconnectivity leading to new cybersecurity challenges. In this publication, we explored the approach of anomaly detection as one possible defense-in-depth solution for protecting the Marine Radar System (MRS). To this end, we evaluated a representative selection of existing approaches and, moreover, presented two novel domain-specific detectors tailored to MRS. Our results show potential for the application of Intrusion Detection Systems (IDSs) for MRS to remedy the current situation. Even if the resilience of MRS network protocols to cyber attacks will be hardened preventively, radar-image-based IDSs could still be beneficial to thwart attacks from the electromagnetic warfare vector. Therefore, the novel detectors we have introduced in this paper promise great value for future systems.

## References

1. Amro, A., et al.: From click to sink: Utilizing ais for command and control in maritime cyber attacks. In: ESORICS (2022). [https://doi.org/10.1007/978-3-031-17143-7\\_26](https://doi.org/10.1007/978-3-031-17143-7_26)

2. Amro, A., et al.: Navigation data anomaly analysis and detection. *Information* **13**(3) (2022). <https://doi.org/10.3390/info13030104>
3. Awan, M.S.K., et al.: Understanding the Vulnerabilities in Digital Components of an Integrated Bridge System (IBS). *JMSE* **7**(10) (2019). <https://doi.org/10.3390/jmse7100350>
4. Balduzzi, M., et al.: A Security Evaluation of AIS Automated Identification System. In: *ACSAC* (2014). <https://doi.org/10.1145/2664243.2664257>
5. Bauer, J., et al.: *Phish & Ships* and Other Delicacies from the Cuisine of Maritime Cyber Attacks. In: *MARESEC* (2023). <https://doi.org/10.5281/zenodo.8406034>
6. Bhatti, J., et al.: Hostile Control of Ships Via False GPS Signals: Demonstration and Detection. *Journal of the Institute of Navigation* **64**(1) (2017). <https://doi.org/10.1002/navi.183>
7. Bole, A., et al.: *Radar and ARPA Manua – Radar and Target Tracking for Professional Mariners, Yachtsmen and Users of Marine Radar*. Butterworth-Heinemann Ltd, 2nd edn. (2009)
8. Caprolu, M., et al.: Vessels Cybersecurity: Issues, Challenges, and the Road Ahead. *IEEE Communications Magazine* **58**(6) (2020). <https://doi.org/10.1109/MCOM.001.1900632>
9. Cohen, S., et al.: Radarnomaly: Protecting radar systems from data manipulation attacks. *Sensors* **22**(11) (2022). <https://doi.org/10.3390/s22114259>
10. Dabrowski, A., et al.: A Digital Interface for Imagery and Control of a Navico/Lowrance Broadband Radar. In: *Proc. of Robotic Sailing* (2011). [https://doi.org/10.1007/978-3-642-22836-0\\_12](https://doi.org/10.1007/978-3-642-22836-0_12)
11. Dagdilelis, D., et al.: Cyber-resilience for marine navigation by information fusion and change detection. *Ocean Engineering* **266** (2022). <https://doi.org/10.1016/j.oceaneng.2022.112605>
12. Hemminghaus, C., et al.: BRAT: a Bridge Attack Tool for Cyber Security Assessments of Maritime Systems. *TransNav* **15**(1) (2021). <https://doi.org/10.12716/1001.15.01.02>
13. IEC 61162-450:2018: Maritime navigation and radiocommunication equipment and systems – Digital interfaces – Part 450: Multiple talkers and multiple listeners – Ethernet interconnection (2018)
14. IEC 61162-460:2018: Maritime navigation and radiocommunication equipment and systems – Digital interfaces – Part 460: Multiple talkers and multiple listeners – Ethernet interconnection – Safety and security (2018)
15. IMO Resolution A.477(12): Performance Standards for Radar Equipment. Resolution, International Maritime Organization (1982)
16. Katsikas, S., et al.: Chapter: Cybersecurity of the Unmanned Ship. *Cybersecurity Issues in Emerging Technologies*, Taylor & Francis Group, 1st edn. (2021)
17. Katsilieris, F., et al.: Detection of malicious ais position spoofing by exploiting radar information. In: *Proc. of the 16th Int. Conf. on Information Fusion* (2013)
18. Kavallieratos, G., et al.: Modelling Shipping 4.0: A Reference Architecture for the Cyber-Enabled Ship. In: *ACIIDS* (2020)
19. Krim Rahaoui, A., et al.: Adaptive threshold for anomaly detection in atm radar data streams. In: *Pattern Recognition and Artificial Intelligence* (2022)
20. Lin, C.Y., et al.: Timing-based anomaly detection in SCADA networks. In: *CRITIS* (2018). [https://doi.org/10.1007/978-3-319-99843-5\\_5](https://doi.org/10.1007/978-3-319-99843-5_5)
21. Longo, G., et al.: Electronic Attacks as a Cyber False Flag against Maritime Radars Systems. In: *IEEE LCN* (2023). <https://doi.org/10.1109/LCN58197.2023.10223370>

22. Longo, G., et al.: Attacking (and Defending) the Maritime Radar System. *IEEE Transactions on Information Forensics and Security* (2023). <https://doi.org/10.1109/TIFS.2023.3282132>
23. Mirsky, Y., et al.: Kitsune: an ensemble of autoencoders for online network intrusion detection. *NDSS* (2018). <https://doi.org/10.14722/ndss.2018.23204>
24. OpenCPN: OpenCPN Chart Plotter. <https://github.com/OpenCPN/OpenCPN> (2023)
25. Packer, J.: Bridge Command. <https://github.com/bridgecommand/bc> (2023)
26. Pavur, J., et al.: A tale of sea and sky on the security of maritime vsat communications. In: *IEEE Symposium on Security and Privacy* (2020). <https://doi.org/10.1109/SP40000.2020.00056>
27. von Rechenberg, M., et al.: Guiding Ship Navigators through the Heavy Seas of Cyberattacks. In: *MARESEC* (2022). <https://doi.org/10.5281/zenodo.7148794>
28. de Riberolles, T., et al.: Characterizing radar network traffic: a first step towards spoofing attack detection. In: *IEEE Aerospace Conference* (2020). <https://doi.org/10.1109/AERO47225.2020.9172292>
29. de Riberolles, T., et al.: Anomaly detection for ICS based on deep learning: a use case for aeronautical radar data. *Annals of Telecommunications* **77**(11) (2022). <https://doi.org/10.1007/s12243-021-00902-7>
30. Roesch, M., et al.: Snort: Lightweight intrusion detection for networks. In: *Lisa*. vol. 99 (1999)
31. SOLAS Chapter V: Safety of Navigation (2009), IMO
32. Spravil, J., et al.: Detecting maritime gps spoofing attacks based on nmea sentence integrity monitoring. *JMSE* **11**(5) (2023). <https://doi.org/10.3390/jmse11050928>
33. Svilicic, B., et al.: Towards a cyber secure shipboard radar. *The Journal of Navigation* **73**(3) (2020). <https://doi.org/10.1017/S0373463319000808>
34. Tam, K., et al.: Macra: a model-based framework for maritime cyber-risk assessment. *WMU Journal of Maritime Affairs* **18** (2019). <https://doi.org/10.1007/s13437-019-00162-2>
35. Wolsing, K., et al.: Network Attacks Against Marine Radar Systems: A Taxonomy, Simulation Environment, and Dataset. In: *IEEE LCN* (2022). <https://doi.org/10.1109/LCN53696.2022.9843801>
36. Wolsing, K., et al.: Can Industrial Intrusion Detection Be SIMPLE? In: *ESORICS* (2022). [https://doi.org/10.1007/978-3-031-17143-7\\_28](https://doi.org/10.1007/978-3-031-17143-7_28)
37. Wolsing, K., et al.: IPAL: Breaking up Silos of Protocol-dependent and Domain-specific Industrial Intrusion Detection Systems. In: *RAID* (2022). <https://doi.org/10.1145/3545948.3545968>