Hostile Control of Ships via False GPS Signals: Demonstration and Detection

JAHSHAN BHATTI and TODD E. HUMPHREYS
The University of Texas at Austin, Austin Texas

Received August 2015; Revised August 2016

ABSTRACT: *An attacker's ability to control a maritime surface vessel by broadcasting counterfeit civil Global Positioning System (GPS) signals is analyzed and demonstrated. The aim of this work is to explore civil maritime transportation's vulnerability to deceptive GPS signals and to develop a detection technique that is compatible with sensors commonly available on modern ships. It is shown that despite access to a variety of high-quality navigation and surveillance sensors, modern maritime navigation depends crucially on satellite navigation and that a deception attack can be disguised as the effects of slowly-changing ocean currents. An innovations-based detection framework that optimally chooses the measurement sampling interval to minimize the probability of a ship exceeding its alert limits without detection is developed and analyzed. A field experiment confirms the vulnerability analysis by demonstrating hostile control of a 65-m yacht in the Mediterranean Sea.*

INTRODUCTION

Surface vessels, from fishing boats to container ships to deep-water oil rigs, depend crucially on Global Positioning System (GPS) signals for navigation, station keeping, and surveillance [1–4]. GPS, its ground and satellite-based augmentation systems, and other Global Navigation Satellite Systems (GNSS) are used as a primary maritime position-fixing system. They are an important maritime navigation aid even for vessels actively piloted by human operators, except in familiar littoral waters such as port entry and within natural or man-made channels where conventional optical navigation is used. Moreover, as surface craft become more autonomous, the trend is toward increased reliance on GNSS: current autopilot systems, dynamic-positioning systems, and fully unmanned surface vehicles are designed under the assumption that GNSS signals are usually available and trustworthy [2, 3, 5, 6]. Even autonomous underwater vehicles typically depend indirectly, or periodically, on GNSS [7].

Given the fragility of GNSS signals under conditions of signal blockage or jamming, and given that the signals do not penetrate underwater, there is interest in developing GNSS-independent maritime navigation and control systems [2, 8]. Terrain-relative navigation has been successfully employed in autonomous submersibles [8], and could serve as a backup to GNSS for surface vessels. This technique has historically required high-resolution (e.g., m-level) underwater terrain maps, which are available for only a tiny fraction of the seafloor, but recent results indicate that coarser (e.g., 20-m-resolution) ship-based bathymetry maps may be adequate for 10-meter-level positioning, provided sufficient terrain variability [9]. Nonetheless, for the present, terrain-relative navigation does not even appear to be an active research topic for civil surface maritime transportation. What is more, the only widespread radionavigation backup to GNSS, Loran-C, was abandoned by the U.S. Coast Guard in 2010 [10], and there are no official U.S. plans for a successor, despite continued lobbying for deployment of its upgrade, eLoran, which is available in other parts of the world [11]. Consequently, one can expect most maritime navigation systems to rely primarily on GNSS for position-fixing for years to come.

By standard practice marine craft are equipped with redundant GNSS units so that one serves as backup if the other experiences a fault. And for extremely critical applications, an entirely GNSS-free positioning system may be available, such as the acoustic positioning system required as a backup to GNSS on dynamically-positioned deepwater drilling vessels [3]. But these fail-safe systems are designed to handle obvious faults or GNSS outages caused by signal blockage or ionospheric effects. They are likely to fail when confronted with a sophisticated and deliberate attacker: outlaws are different from outliers; fraud is different from faults.

NAVIGATION: *Journal of The Institute of Navigation*
Vol. 64, No. 1, Spring 2017
Printed in the U.S.A.

A GNSS deception attack, in which counterfeit GNSS signals are generated for the purpose of manipulating a target receiver's reported position, velocity, or time, is a potentially dangerous tool in the hands of a deliberate attacker. While there have been no confirmed reports of such attacks performed with malice, convincing demonstrations have been conducted both in the laboratory and in the field with low-cost equipment against a wide variety of GPS receivers [12–14]. The key to the success of these so-called GPS spoofing attacks is that, whereas the military GPS waveforms are by design unpredictable and therefore resistant to spoofing, civil GPS waveforms—and those of other civil GNSS—are unencrypted, unauthenticated, and openly specified in publicly-available documents [15, 16]. Also, although not entirely constrained by the GNSS signal specifications, the navigation data messages modulating these civil waveforms are highly predictable. The combination of known signal structure and navigation data predictability makes civil GNSS signals an easy target for spoofing attacks.

The departure point for development of a spoofing detection framework is the impressive corpus of fault detection and isolation (FDI) literature, the result of more than four decades of effort. Sensor deception can be thought of as a special type of sensor fault in which a strategic attacker has some level of control over the fault behavior and applies this control with malicious intent. Several classes of methods for sensor FDI in stochastic linear dynamic systems are surveyed in [17–20]. Although many sophisticated approaches have been developed in this mature field, most fault-detection methods focus on minimizing time-to-detect without regard to integrity risk, as noted by Joerger [21]. Integrity risk is the appropriate figure of merit for dynamic systems with clearly specified alert limits such as are common in aviation and maritime navigation and in time transfer. For these systems, state estimation errors that remain within the alert limits cause no performance degradation or heightened safety risk, but undetected errors exceeding the alert limit can have severe consequences.

The first attempt to address sensor deception by minimizing integrity risk appears to be [22], where a model-based spoofing detection method was developed for an aircraft's GPS-aided inertial navigation system. However, the analysis considered a batch detection test whose batch interval is aligned with the attack interval, a coincidence that cannot be expected in practice. The current work adopts a sequential detection approach, which is more appropriate for attacks of unknown start time and duration. But as opposed to sequential detection techniques designed to minimize time-to-detect for fixed probabilities of false alarm and detection, such as the sequential probability ratio test [23], the current work adopts a fixed time-to-detect approach and follows [21] and [22] in seeking to minimize integrity risk. More precisely, this work minimizes mean integrity risk, or integrity risk averaged over all possible attack start times.

The heart of a detection technique is the so-called detection statistic, a function of the sensor measurements that gets compared to a threshold [24]. This work adopts an innovations-based detection statistic whose performance is insensitive to the particular time history of false differential position and velocity induced by the attacker.

A key feature of the current work's detection framework is that it optimizes the measurement sampling interval; the standard innovations-based detection approach makes no attempt at such optimization [17, 25]. The optimization seeks toz minimize worst-case integrity risk over a set of reasonable attack profiles. Measurement sampling interval optimization was previously considered in [26], but that work minimized time-to-detect whereas the current work's criterion is integrity risk.

This paper makes three contributions. First, it details the pathways and effects of GNSS deception on maritime navigation and surveillance. Whereas maritime transportation's vulnerability to GNSS jamming has been previously established [2], this work offers the first detailed analysis of the effects of GNSS deception on a surface vessel. Second, it develops an innovations-based spoofing detection framework and optimizes the worst-case mean integrity risk within this framework given a set of reasonable attack profiles. Third, it presents the results of an unprecedented field experiment demonstrating hostile control of a 65-m yacht in the Mediterranean Sea.

GNSS DEPENDENCIES OF A MODERN INTEGRATED BRIDGE SYSTEM

This section details the pathways and effects of GNSS deception on maritime navigation and surveillance. Besides providing a deeper understanding of the vulnerability of maritime vessels to GNSS spoofing, this overview will identify a subset of ship sensors that can conveniently and effectively be applied to the problem of spoofing detection.

Compass

The magnetic cand gyrocompass (a gyroscope designed to be north-seeking by taking advantage of Earth's rotation) depend only weakly on GNSS. A magnetic compass requires knowledge of latitude and longitude to correct for magnetic variation [27]. A gyrocompass requires knowledge of latitude and speed in the north/south direction to correct for

"northing" error [27]. However, outside of the polar regions, position errors on the order of tens of kilometers and velocity errors on the order of meters per second will only cause pointing errors on the order of a degree. Therefore, this work will neither exploit nor model the weak coupling between GNSS and traditional ship compasses.

However, a satellite compass [28], which provides both the position and three-axis attitude of the ship, is fully reliant on GNSS. A common satellite compass comprises two GNSS receivers separated by a 0.2–10 meter baseline coupled with miniature accelerometers, gyros, and a magnetometer. The low cost, size, weight, and power consumption of satellite compasses, and the fact that they never require calibration, make these devices an increasingly popular compass option for surface vessels.

Collision Avoidance

The Automatic Radar Plotting Aid (ARPA) and the view from the bridge windows are the primary means used for collision avoidance. The ARPA processes and displays the raw radar data in a polar azimuth-range plot, tracks targets, and computes time and distance of closest approach for each target [29]. Without the additional information that sensors like compass, speed log, and GNSS provide, the ARPA can still perform collision-avoidance functions but can only display target information oriented along the ship's heading, the so-called heads-up mode, with relative motion. With compass information, the ARPA can present the radar data oriented along the ship's velocity vector, the so-called course-up mode, which prevents smearing of the returns during course-change maneuvers. Similarly, the ARPA can present the radar data in a so-called true motion mode, where the motion is either sea-stabilized by compass and speed log or ground-stabilized by GNSS. (An interesting effect of a GNSS deception attack with ground stabilization enabled on the ARPA is that radar echos from land masses appear to move when they should be stationary.) GNSS information also allows the ARPA to compute latitude and longitude for the tracked targets. Nevertheless, convenience features such as ground stabilization and target localization that depend on GNSS signals play a relatively minor role in collision avoidance with other moving targets.

The Automatic Identification System (AIS) allows ships to communicate their position, heading, and speed in a self-organizing radio network to aid in collision avoidance [27]. A ship's AIS transceiver typically relies on a GNSS-based positioning source, although it can revert to a pre-determined backup source during a manifest GNSS failure. Under a GNSS deception attack, a ship may transmit misleading AIS reports and incorrectly compute the point of closest approach to surrounding ships, raising the collision risk. An ARPA can typically overlay the AIS over the radar return, and a modern ARPA with integrated AIS can automatically correlate AIS and radar positions into a single target.

Dead Reckoning

Dead reckoning (DR) is the process of propagating a known position based solely on a ship's course and speed, derived from compass and speed log measurements. An estimated position (EP) corrects a dead-reckoned position by applying approximate knowledge of the effects of environmental disturbances such as leeway (drift due to wind), and tidal and ocean currents. Typically, the effects of environmental disturbances are lumped together into a velocity error vector, whose angle and magnitude are referred to as set and drift, respectively. The set and drift can be estimated by comparing a dead-reckoned position to a position fix derived from either a GNSS receiver (typically), observations of celestial bodies, or radar and visual bearings. On paper charts, DR would be reset with a position fix at least every hour, or as often as every three minutes, depending on the accuracy required for navigating the surrounding waters [27]. Electronic chart systems, discussed in the next section, all have the ability to automate DR, making it easier to detect GNSS faults or deception.

Electronic Chart Display and Information System

The Electronic Chart Display and Information System (ECDIS) consolidates the measurements available from various ship sensors and integrates systems such as ARPA, AIS, and DR as shown in Figure 1 to provide complete situational awareness to the ship's crew [27]. ECDIS is the primary tool for route planning and tertiary to the ARPA and AIS for collision avoidance, as mandated by legislation and made explicit in maritime training. Most ECDIS allow overlaying ARPA and AIS information on the charts and planned route for convenience. The overlay may be useful in detecting discrepancies that would arise due to GNSS deception of the own-ship position, e.g., failure of radar returns to match coastal features and buoys on the charts, or the AIS-reported position of nearby vessels. Maritime training emphasizes the need to look for and investigate discrepancies as they normally indicate an equipment problem. But these discrepancies may simply confuse a crew unaware of GNSS deception, despite training manuals that have begun to identify GNSS deception as a potential threat [30]. In any case, when the distance to shore exceeds the range of radar (20 km for low-frequency radar, less for X-band) and when there are few ships nearby,

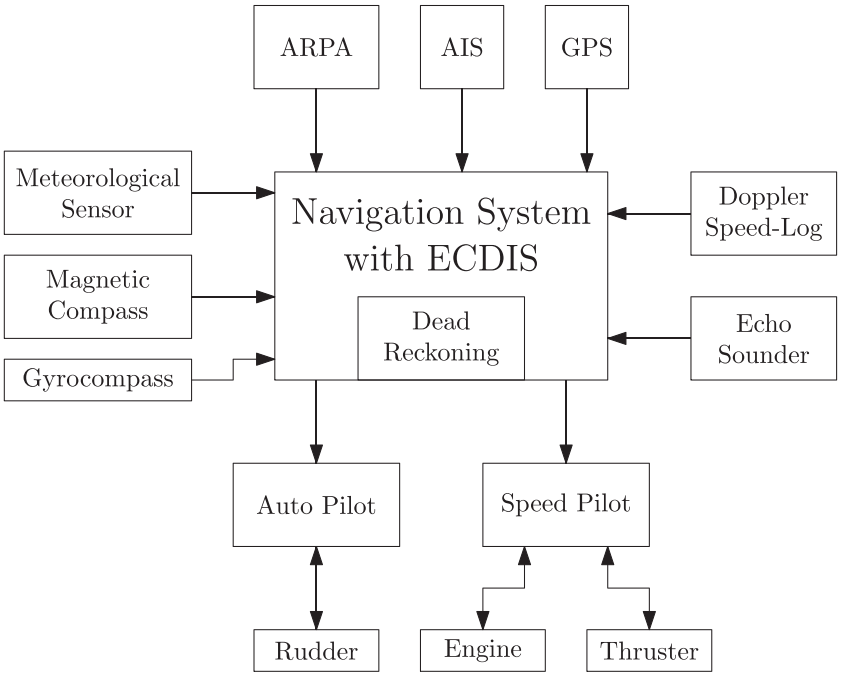

Fig. 1–Block diagram showing relationship between sensors, actuators, and the ECDIS on a modern integrated bridge system.

GNSS deception attacks are not likely to be detected solely with radar. Most electronic chart systems such as the Totem ECDIS allow configuring the reset interval of the built-in DR and raising an alarm if the difference between the position fix and DR exceeds a threshold [31]. Section IV will develop an analytically rigorous foundation for this approach by relating the detection threshold to the probability of hazardously misleading information (HMI) for a given false alarm rate and fix interval.

Autopilot System

Virtually, all large ships have a course autopilot, which maintains a prescribed heading through rudder actuation in response to compass feedback. Some ships will additionally have a speed autopilot, which maintains a prescribed speed through water by varying the engine thrust in response to feedback from the Doppler speed log sensor. Neither of these rudimentary autopilot systems depends on GNSS directly. However, the course autopilot is typically driven by a higher-level track-keeping system that requires GNSS feedback. This work focuses on conventional PID-based control systems because they are commonly implemented in practice and typically perform just as well as adaptive model-based control systems under nominal sea and ship conditions [32].

GNSS-Independent Sensors

Sensors which do not have any dependency on GNSS include inertial, acoustic, visual, and meteorological sensors. An inertial sensor found on most ships is a gyroscope-based rate-of-turn (ROT) sensor, which is independent of the compass and GNSS, for derivative course control feedback. The modern speed log uses acoustic Doppler measurements from particles in the water column to compute three-axis speed through water. Other acoustic sensors include conventional downward-looking sonar, also known as an echo sounder, for sea depth measurements and round-trip acoustic ranging to transponders embedded in the sea floor for dynamic positioning [3]. Meteorological sensors provide measurements of temperature, wind, and pressure that can help predict, for example, the effect of leeway and surface currents [27]. Visual bearing measurements of known reference points such as terrestrial landmarks or celestial bodies can be used for positioning. Celestial navigation requires knowing the time, either from GNSS or a free-running quartz crystal clock, to look up the position of celestial bodies from an almanac [27]. A jump in ship time by 5 sec (e.g., due to leap second spoofing) would cause a longitude error of 0.02 deg. Nevertheless, errors less than 10 sec from either a drifting or GNSS-deceived clock are comparable with other errors in celestial navigation.

These GNSS-independent sensors feed into alternative position sources that could be used to cross-check GNSS in a modern integrated bridge system. However, a subtle-enough spoofing attack can be consistent with dead reckoning or celestial navigation and thus escape detection. Also, acoustic positioning is only useful for vessels operating in the small neighborhood of the transponders. Although

this work's focus is on GNSS deception, it is worth mentioning that radar and acoustic sensor systems on modern civil surface vessels are also vulnerable to deception and jamming. Thus, although these systems are assumed herein to be trustworthy and potentially useful for detecting GNSS deception, a more thorough security analysis would need to consider a coordinated, self-consistent attack on GNSS, radar, and acoustic sensors.

Summary of GNSS Deception Vulnerabilities

The ship's crew can cross-check GNSS with the (1) compass, (2) speed log, (3) ship dynamics model, (4) radar returns compared with AIS from other ships, (5) radar returns off buoys and coastlines compared with charts, (6) echo sounder, and (7) meteorological sensors. But, as will be shown later on, even an optimal combination of (1)–(3), which amounts to sophisticated DR, would not be sufficient to reliably detect a subtle attack before the ship's positioning error exceeds a reasonable hazardous condition threshold. If (4) and (5) are properly and fully exploited, then the security situation improves significantly. But alignment of charted objects such as the shoreline and buoys with radar returns is often quite poor even under normal conditions because of (i) shoreline changes with tide, (ii) inadequate resolution of charts, and (iii) positioning, bearing, and radar-ranging errors. Consequently many ships' crews either do not attempt radar overlay or would not consider it a trustworthy cross-check for ownship positioning errors. Also, comparing radar with AIS from other ships is not trustworthy because AIS data can be easily manipulated and AIS-repeated location data ultimately depend on GNSS.

For avoidance of collisions with radar-reflective objects, the ARPA remains trustworthy and its collision avoidance function does not depend on GNSS. But cross-track ship excursions outside the planned corridor are nevertheless dangerous precisely because some threatening objects (e.g., underwater hazards) are not visible to radar and will not be detected by downward-looking sonar. Moreover, along-track errors in a ship's position can also be hazardous because such errors can confuse the interpretation of radar returns or cause a ship to over- or under-shoot the location of a planned maneuver.

Illustrative Example: The Grounding of the Royal Majesty

For us to appreciate the possible effects of a GNSS deception attack on a surface vessel, it is instructive to consider the grounding of the 174-m cruise liner Royal Majesty [33, 34]. Shortly after the ship departed Bermuda for Boston in June of 1995, the cable connecting its GPS antenna to the unit on the bridge became detached, forcing the GPS unit to transition to a DR mode in which the ship's location was extrapolated from the last known good location based solely on gyro compass and water speed measurements. The crew and autopilot, unaware of the transition to DR mode, accepted the position indicated on the radar display's map as truthful even as the ship accumulated a 31 km cross-track navigational error. As the ship approached Nantucket, the crew misidentified one buoy and ignored the absence of another. The ship's GPS-based navigation system had performed so utterly reliably in the past that the crew's trust in the ship's displayed position was not shaken even as a lookout sighted blue and white water ahead. Minutes later, the ship ran aground on shoals invisible to its radar and sonar system.

In the aftermath of the Royal Majesty grounding, integrated bridge systems were modified to more clearly indicate loss of GNSS signals, and redundant GNSS units became standard. In addition, the incident is used as an important lesson on the dangers of over-reliance on GNSS in maritime training colleges since the crew of the Royal Majesty clearly acted in a manner inconsistent with proper training, a contributing cause to the incident. Nevertheless, the risk of a repeat of the Royal Majesty grounding, or a similar incident, caused by deliberate, strategic GNSS deception remains because there would be no apparent loss of GNSS, the DR would appear to remain consistent with GNSS, and because primary and backup GNSS units would be equivalently affected.

Having offered an overview of the possible effects of GNSS deception on surface vessels, this paper now turns to developing a framework for analysis of GNSS spoofing detection based on comparison of GNSS data with a modified version of DR. This detection strategy is appealing because of its broad applicability: all sizable surface vessels can perform at least rudimentary DR, and the DR technique works both far from shore and in littoral waters. The next two sections introduce the dynamics model and the detection framework.

SHIP AND SPOOFING MODEL

Consider a simplified ship dynamics model with a conventional track-keeping guidance system. A conventional track-keeping system attempts to zero the ship's cross-track position using a proportional-integral (PI) controller wrapped around a course autopilot, as shown in Figure 2.

Ship Dynamics

The ship dynamics model presented here, although simple compared to a more expressive

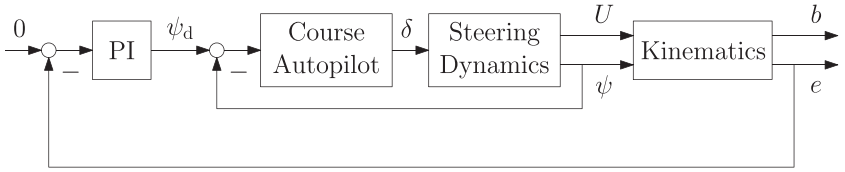

Fig. 2–Conventional track-keeping system based on an existing course autopilot system [32, p. 293]. *Here, ψ_d is the desired heading angle, δ is the rudder angle, U is the ship speed through water, ψ is the heading angle, b is the along-track position, and e is the cross-track position.*

six degree-of-freedom model, captures the low-frequency ship motion relevant for control and spoofing. The ship's steering dynamics is described by a first order Nomoto model [32],

$$T\dot{r} + r = K\delta + r_b,$$

where T is the ship's time constant (s), K is the rudder gain (1/s), δ is the rudder angle (rad), r is the ship's turn rate (rad/s), and r_b is a slowly-varying parameter that models environmental disturbances (rad/s). The rudder angle δ and angular rate $\dot{\delta}$ are physically constrained by saturation conditions $|\delta| < \delta_{max}$ and $|\dot{\delta}| < \dot{\delta}_{max}$, respectively, but the controller is designed such that the rudder angle dynamics remain linear under typical conditions. The ship's kinematics are given by [32]

$$\begin{aligned}\dot{\psi} &= r\\ \dot{x} &= U\cos\psi + d_x\\ \dot{y} &= U\sin\psi + d_y,\end{aligned}$$

where U is the ship's speed through water (m/s), d_x and d_y model errors due to drift caused by slowly-varying environmental disturbances such as ocean currents and wind (m/s), x and y are the ship's northing and easting (m), respectively, and ψ is the ship's heading (rad). Zero heading is defined to be due north with increasing heading clockwise. The environmental disturbance parameters are modeled as Gauss-Markov processes,

$$\begin{aligned}\dot{d}_x &= -\frac{1}{T_d}d_x + v_x\\ \dot{d}_y &= -\frac{1}{T_d}d_y + v_y,\end{aligned}$$

where T_d is the disturbance time constant and v_x and v_y are additive white Gaussian noise (AWGN) sources with intensity σ_d^2 (m^2/s^3).

Ship Control Laws

Only conventional controllers are considered in the sequel because they perform just as well as adaptive and non-linear model-based controllers under nominal sea and ship conditions [32]. A conventional course autopilot controls the ship's heading ψ to a desired approximately-constant heading ψ_d using a proportional-integral-derivative (PID) control law. In modeling the course control law that follows, and the track-keeping control law presented thereafter, the measurements are assumed to be noiseless and continuous since the low-bandwidth controllers and ship dynamics act to suppress the effects of real-world discretization and measurement noise at the output of each closed-loop control system. The measurements $\psi(t)$ and $r(t)$ from the compass and ROT sensor, respectively, control the rudder angle $\delta(t)$ according to

$$\delta(t) = K_i\int_0^t \left[\psi_d - \psi(\tau)\right]d\tau + K_p\left[\psi_d - \psi(t)\right] - K_d r(t),$$

where K_i is the integral gain, K_p is the proportional gain, and K_d is the derivative gain. Following conventional PID control design of second-order systems [32, p. 261], these gain parameters are derived from a chosen natural frequency ω_n and relative damping ratio ξ of the closed-loop system; the latter is typically chosen in the interval $0.8 \le \xi \le 1.0$. The closed-loop bandwidth, ω_b, defined as

$$\omega_b \triangleq \omega_n\sqrt{1 - 2\xi^2 + \sqrt{4\xi^4 - 4\xi^2 + 2}},$$

is chosen such that

$$\frac{1}{T} < \omega_b < \omega_\delta,$$

where $\omega_\delta \triangleq \frac{\dot{\delta}_{max}}{\delta_{max}}$ is the rudder servo bandwidth. Finally, the PID gains are related to ω_n and ξ by

$$\begin{aligned}K_p &= \frac{T}{K}\omega_n^2\\ K_d &= \frac{1}{K}\left[2T\xi\omega_n - 1\right]\\ K_i &= \frac{T}{K}\frac{\omega_n^3}{10}.\end{aligned}$$

An outer control loop for track-keeping is typically wrapped around the course autopilot. In some cases, a human operator in the loop may take the role of track-keeping controller. Whether mechanical or human, the controller can be modeled as a PI controller. The track, or rhumb line, can be approximated in the local Cartesian coordinates by a ray, which is parametrized by an angle ψ_0 (rad) and start position x_0 and y_0 (m). The along-track and cross-track position, b and e, respectively, are given

by

$$b = (x - x_0)\cos\psi_0 + (y - y_0)\sin\psi_0$$
$$e = (y - y_0)\cos\psi_0 - (x - x_0)\sin\psi_0.$$

The relationship between the global and track coordinates is illustrated graphically in Figure 3.

Because GNSS is the most accurate positioning source, nearly always available, and assumed to be reliable when available, it is typically the primary positioning source [2]. The GNSS receiver's cross-track position measurement, which is taken to be equivalent to $e(t)$, is fed back with a PI control law given by

$$\psi_d(t) = \psi_0 - K_i' \int_0^t e(\tau)d\tau - K_p' e(t),$$

where K_i' is the integral gain and K_p' is the proportional gain. The gains are chosen so that the inner course control loop and the outer track-keeping loop have significant time scale separation, with the inner loop faster, a typical practice for marine and aerial vehicle cascaded controller design [32, 35]. Thus, from the perspective of the outer loop, one can assume $\psi \approx \psi_d$, and the full closed-loop cross-track dynamics can be approximated by a first-order system with bandwidth $\omega_b' = UK_p' \ll \omega_b$. Note that the along-track position is not controlled by a feedback law but instead proceeds open-loop according to an approximate crew-selected velocity setpoint.

Spoofer Control Law

In a spoofing attack, the ship's GNSS receiver will report the position commanded by the spoofer. To remain covert, the spoofer will typically command positions that are gentle deviations, conveniently represented in along-track and cross-track coordinates, from the ship's true position. Cross-track deviations will prompt a response from the ship's track-keeping controller whereas along-track deviations will elicit no response unless the ship's track changes. Along-track spoofing can be an effective strategy from the point of view of the attacker, but this paper will focus on cross-track spoofing because it is equally effective yet requires less knowledge of the ship's route.

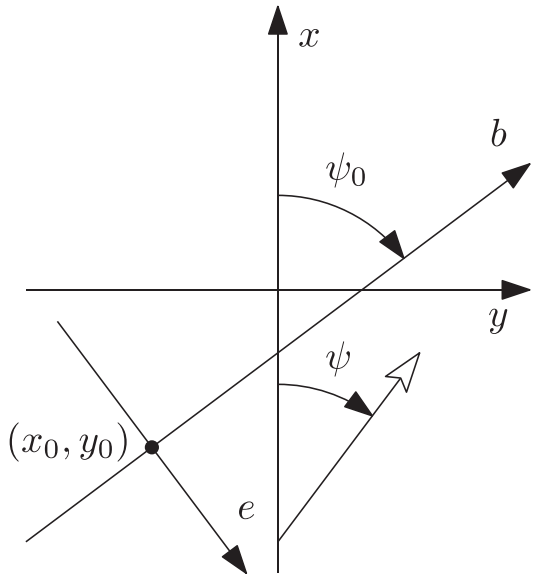

Fig. 3–Coordinate systems for ship global position (x,y) and track position (b,e). The track coordinate system's origin and rotation with respect to the global coordinate system is given by (x_0,y_0) and ψ_0, respectively. The ship's orientation with respect to the global coordinate system is given by heading angle ψ.

In a cross-track spoofing attack, the spoofer generates a GNSS signal whose implied coordinates are the attacker's estimate of the ship's actual along-track position $\hat{b}_a(t)$ and a spoofed cross-track position $e_s(t)$. The latter can be written as the difference of two parts, the attacker's estimate of the ship's true cross-track position $\hat{e}_a(t)$ and a spoofer-induced cross-track modulation $e_m(t)$ so that $e_s(t) = \hat{e}_a(t) - e_m(t)$. The modulation $e_m(t)$ is also called the attack profile, as it represents the spoofer-intended departure from the cross-track position. Note that to form $\hat{e}_a(t)$, the attacker must continuously estimate both the ship's position and its rhumb line. This assumption is not particularly demanding: other-ship position estimation via radar is both accurate and routine, and surface vessels typically follow a route consisting of waypoints connected by readily-estimable lines of constant bearing.

The attacker's goal is to force the ship to track a spoofer-commanded cross-track position, denoted $\bar{e}$, as quickly as possible without being detected. He evades detection by generating a subtle $e_m(t)$ with limited velocity and acceleration magnitudes:

$$|\dot{e}_m(t)| \le v_{max}, \quad |\ddot{e}_m(t)| \le u_{max} \tag{1}$$

Solving the following minimum-time optimal-control problem yields the attack profile $e_m(t)$ that achieves the spoofer's goal. Here, t_f is the final time, and the control input $u(t)$ enters through the second derivative of $e_m(t)$ as part of the dynamic constraint:

$$\begin{aligned} \min_{u(t)} \quad & t_f \\ \text{s. t.} \quad & \ddot{e}_m(t) = u(t) \\ & e_m(0) = 0,\ \dot{e}_m(0) = 0 \\ & e_m(t_f) = \bar{e},\ \dot{e}_m(t_f) = 0 \\ & |\dot{e}_m(t)| \le v_{max} \\ & |\ddot{e}_m(t)| \le u_{max}. \end{aligned} \tag{2}$$

For $\bar{e} \to \infty$ and $t_c \triangleq \frac{v_{max}}{u_{max}}$, the solution is given by

$$e_m(t) = \begin{cases} \frac{1}{2}u_{max}t^2 & 0 < t \le t_c \\ \frac{1}{2}u_{max}t_c^2 + v_{max}(t - t_c) & t_c < t \end{cases}$$

An attack profile generated as a solution to (2) is easily disguised as the effect of ocean currents. But it may not be optimal from the point of view of the attacker; i.e., it may not be the most hazardous undetectable profile. The optimal profile in this sense actually depends on the defender's particular detection test. Strategies for generating $e_{\mathrm{m}}(t)$ that are more directly related to plausible detection tests are considered in [36, Appendix A]. Nonetheless, the strategy outlined in (2) has the virtue of being intuitive and readily implementable yet generates $e_{\mathrm{m}}(t)$ profiles similar to those produced by the more complex strategies.

The maxima $v_{\max}$ and $u_{\max}$ are assumed to be sufficiently small that $e_{\mathrm{m}}(t)$ is slow compared to the time constant of the ship's track-keeping control law, ensuring the attacker can dictate the ship's true cross-track position $e(t)$ with only modest errors—errors due to the spoofer's imperfect estimate of the ship's true position and of the rhumb line, and to the ship's own estimation and control errors. Under this assumption, the spoofer needs not adapt $e_{\mathrm{m}}(t)$ to the ship's response but may simply generate $e_{\mathrm{m}}(t)$ open loop. A closed-loop spoofing controller is also possible, but its attacks are more difficult to maintain covert, as illustrated in [14].

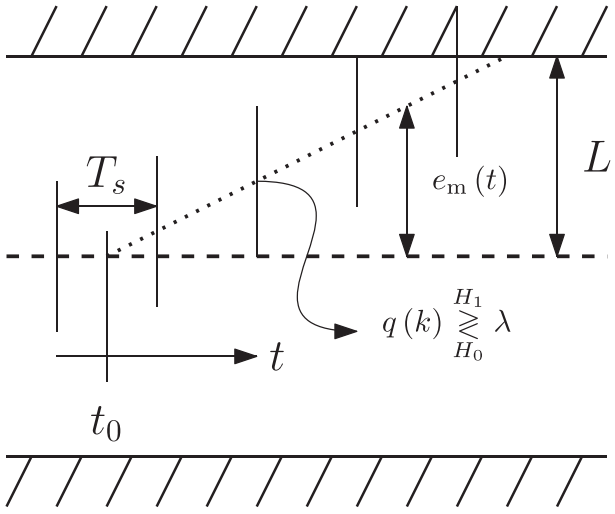

Fig. 4–A graphical overview of the detection problem: A spoofing attack with cross-track profile $e_{\mathrm{m}}(t)$ begins at time t_0, which is unknown to the defender. The attacker attempts to drive the ship to exceed the alert limit L, beyond which lie potential hazards, without detection. At every time $t_k = kT_s$, $k = 0, 1, ...$, a GNSS measurement is taken and used to form the detection statistic $q(k)$. The time instants t_k are unknown to the attacker, though the measurement period T_s may be known. If $q(k)$ exceeds the threshold λ, the alternative hypothesis H_1 (spoofing attack) is declared; otherwise, the null hypothesis H_0 is assumed.

DETECTION FRAMEWORK

The detection framework developed in this paper attempts to minimize the mean integrity risk $\bar{I}_R$, defined subsequently, for a given continuity risk $C_R \triangleq 1/M_F$, where M_F is the mean time between false alarms. This framework borrows concepts from GNSS integrity monitoring in aviation applications [21] and the fault detection literature [17], which are applied here to the "fraud detection" problem. The introduction of mean integrity risk, which is a marginal probability, departs from the usual definition of integrity risk as adopted by the aviation community. However, it is a necessary adaptation to account for an uncoordinated spoofer and defender, as will be shown clearly in the sequel. Typically, the integrity and continuity risk are specified in terms of the probability of hazardously misleading information (HMI) per approach and the false-alarm rate, respectively.

Overview

The schematic in Figure 4 offers a graphic overview of the detection problem. Time $t = 0$ denotes the beginning of an approach, or part of a journey, such as the final approach to a harbor. At each time $t_k = kT_s$, $k = 0, 1, ...$, a detection test is performed to decide between two hypotheses—the null hypothesis H_0 indicating nominal operating conditions, and the alternative hypothesis H_1 indicating a spoofing attack is underway. At the beginning of the approach, H_0 is true; at some time $t_0 \geq 0$, a transition to H_1 occurs. After t_0, the attack continues until either hazardous conditions occur or the attack is detected. In this framework, the constant time between tests T_s is a key parameter: it is taken as the free parameter for the integrity optimization problem.

The detection strategy envisioned here is decoupled from the ship's track-keeping controller, which is assumed to ingest GNSS measurements at its usual rate—typically much faster than $1/T_s$—without regard for the the periodic detection tests occurring in parallel. A joint control-and-detection framework is possible and would have slightly superior performance compared to the proposed framework, but a disjoint framework is simpler and has the benefit of being applicable to existing ships without re-certification of their integrated bridge systems.

So long as the detection statistic $q(k)$ remains below a threshold λ, the detector assumes H_0; otherwise, it declares H_1 and continuity is broken as the crew attempt to neutralize the potential spoofing threat. The threshold λ is chosen to satisfy

$$P\left(q(k) > \lambda \,|\, H_0\right) = T_s/M_F = C_R T_s$$

to maintain the prescribed false-alarm rate. Note that it will be shown that the probability distribution of $q(k)$ under H_0 is independent of k, so λ need not depend on k.

Integrity Risk

Leading up to a definition of mean integrity risk $\bar{I}_R$, it will be useful to define what is meant by hazardous conditions and by a so-called local HMI event. Let the total system error of a certain state element of interest be denoted $\epsilon(t)$. The total system error is the departure of the true state element from the controller's desired value of that state element, and includes both estimation and control errors. Hazardous conditions are said to occur when $|\epsilon| > L$ for an alert limit $L > 0$. Although the ship may not be in immediate danger if $|\epsilon| > L$, control decisions based on such divergent estimates are highly risky. In this paper, the state of interest is the cross-track position $e(t)$, and a typical value for L may be 1 km. To account for worst-case control error, L must be substantially smaller than the distance that the ship's route clears charted hazards.

Assuming GNSS measurements are continuously available, as in the ship's control model, and that control errors remain small, then under H_1, $\epsilon(t) \approx e_{\mathrm{m}}(t)$. This deterministic approximation is a key simplifying assumption: it prevents the total system error from being correlated with the detection statistic. Lack of correlation greatly simplifies the expression for the mean integrity risk, as will be shown subsequently.

A local HMI event $E(t)$ for $t > t_0$ is defined as hazardous conditions under a spoofing attack that has not been detected. Mathematically, $E(t)$ is expressed as

$$E(t) \triangleq \left(|e_{\mathrm{m}}(t)| > L\right) \wedge \left(\bigwedge_{k \in S_t} q(k) < \lambda\right),$$

where $S_t \triangleq \{k \,|\, t_0 < kT_s < t\}$. The boolean event G indicates whether a local HMI event has occurred at any time $t > t_0$ during an approach:

$$G \triangleq \bigvee_{t > t_0} E(t).$$

Let the first time hazardous conditions occur under a spoofing attack be denoted t_L and let $S_L \triangleq \{k \,|\, t_0 < kT_s \leq t_L\}$. Then G can be reformulated as

$$G = \bigwedge_{k \in S_L} q(k) < \lambda.$$

Integrity risk is defined for a particular start time t_0 as

$$I_R(t_0) \triangleq P\left(G \,|\, H_0\right) P\left(H_0\right) + P\left(G \,|\, H_1, t_0\right) P\left(H_1\right),$$

where $P(H_0)$ and $P(H_1) = 1 - P(H_0)$ are the prior probabilities for H_0 and H_1.

In this work, the integrity risk is assumed to be dominated by $P\left(G \,|\, H_1, t_0\right)$ for the purposes of the optimization problem defined in the sequel since we are maximizing the integrity risk with respect to the spoofer attack parameters. An attack with $P\left(G \,|\, H_1, t_0\right) < P\left(G \,|\, H_0\right)$ is not particularly effective, so the probability of a spoofing attack $P\left(H_1\right)$ is conservatively assumed to be unity. Of course, care must be taken to verify that the inequality condition is actually true, which is ensured by a sufficiently large value of L.

In aviation applications, the integrity risk is typically accompanied by a time-to-alarm requirement. In the previous formulation, the time-to-alarm t_A is implicitly zero seconds, which is the strictest possible requirement. However, a non-zero t_A can be accounted for by simply modifying

$$S_L \triangleq \left\{k \,|\, t_0 < kT_s \leq t_L + t_A\right\}.$$

Taking all spoofing start times to be equally likely and assuming the attack is time-invariant, the mean integrity risk can be defined as

$$\bar{I}_R \triangleq \int_0^1 P(G \,|\, H_1, t_0 = \beta T_s)\, d\beta. \tag{3}$$

The time-invariance property stems from the deterministic approximation of $\epsilon(t)$ made in the beginning of this section so that for any non-negative integers l, m and $0 \leq \beta \leq 1$,

$$P(G \,|\, H_1, t_0 = \left(l + \beta\right) T_s) = P(G \,|\, H_1, t_0 = (m + \beta)\, T_s).$$

Detection Statistic

Detector performance depends strongly on the detection statistic $q(k)$. If the attack profile $e_{\mathrm{m}}(t)$ were precisely known to the defender *a priori*, then a detection statistic could be optimally tailored to the known profile. The statistic would amount to processing estimator innovations through a filter matched to the known profile [36, Appendix A]. If some attack profile parameters remained unknown, such as t_0 and $v_{\max}$, then the generalized likelihood ratio approach would be reasonable [17]. However, the stronger the defender's assumptions are about the attack profile, the more vulnerable he becomes to an attacker who violates those assumptions.

One recognizes a zero-sum game in the simultaneous incentive: the defender has to optimize $q(k)$ for the attacker's choice of $e_{\mathrm{m}}(t)$ and the attacker has to optimize $e_{\mathrm{m}}(t)$ for the defender's choice of $q(k)$. If an equilibrium pair $\{q^\star(k), e_{\mathrm{m}}^\star(t)\}$ were found to exist for this game, such that neither attacker nor defender benefits by unilateral departure from the equilibrium, then $q^\star(k)$ could be taken as an equilibrium-optimal detection statistic [37, 38]. However, the authors were unable to discover such an equilibrium; its existence remains an open question. Instead, a normalized-innovations-squared (NIS) statistic [17, 25, 39] is adopted here.

This statistic is not optimal in the sense of $q^\star(k)$ but is robust in that it makes no assumptions about the attack trajectory; rather, it penalizes all departures from the assumed model.

The innovation sequence $\nu(k)$ on which $q(k)$ is based is generated by a Kalman filter ingesting GNSS measurements every T_s seconds. A simplified model for the Kalman filter is developed below in preparation for determining the probability distribution of $q(k)$. First, consider the continuous-time ship dynamics model

$$\dot{\eta}(t) = A\eta(t) + Bu(t) + \Gamma\tilde{v}(t),$$

where

$\eta = \begin{bmatrix} x & y & d_x & d_y \end{bmatrix}^{\mathrm{T}}$ is the state vector,

$A = \begin{bmatrix} 0 & I_2 \\ 0 & -\frac{1}{T_d} I_2 \end{bmatrix}, B = \begin{bmatrix} I_2 \\ 0 \end{bmatrix}, \Gamma = \begin{bmatrix} 0 \\ I_2 \end{bmatrix},$

$u = U\begin{bmatrix} \sin\psi & \cos\psi \end{bmatrix}^{\mathrm{T}}$ is the control, and

$\tilde{v} = \begin{bmatrix} v_x & v_y \end{bmatrix}^{\mathrm{T}}$ is AWGN with intensity $Q_c = \sigma_d^2 I_2$,

with I_n the n-by-n identity matrix and 0 an appropriately-dimensioned zero matrix. The control $u(t)$ is derived from the ship's compass and speed log measurements. The potentially-spoofed GNSS measurements are sampled from

$$z(k) = H\eta(kT_s) - z_{\mathrm{m}}(kT_s) + w(k),$$

where $w(k)$ is a discrete AWGN sequence with covariance $R = \sigma_p^2 I_2$, $H = \begin{bmatrix} I_2 & 0 \end{bmatrix}$, and $z_{\mathrm{m}}(t)$ is the deterministic spoofer-induced two-dimensional position modulation for which, by definition, $z_{\mathrm{m}}(t) = 0$ for $t < t_0$.

The *a priori* and *a posteriori* Kalman filter estimates $\bar{\eta}(k)$ and $\hat{\eta}(k)$ are related to the corresponding estimation errors by $\bar{\epsilon}(k) \triangleq \eta(k) - \bar{\eta}(k)$ and $\hat{\epsilon}(k) \triangleq \eta(k) - \hat{\eta}(k)$. The filter innovation $\nu(k) \triangleq z(k) - H\bar{\eta}(k)$ is equivalent to

$$\nu(k) = H\bar{\epsilon}(k) - z_{\mathrm{m}}(kT_s) + w(k).$$

The recursion equations for the estimation error's means and covariances are given by

$$\begin{aligned}
\mathbb{E}\left[\bar{\epsilon}(k)\right] &= F\mathbb{E}\left[\hat{\epsilon}(k-1)\right] \\
\bar{P}(k) &\triangleq \mathbb{E}\left[\bar{\epsilon}(k)\bar{\epsilon}^T(k)\right] = FP(k-1)F^{\mathrm{T}} + Q \\
\mathbb{E}\left[\hat{\epsilon}(k)\right] &= \left(I - K(k)H\right)\mathbb{E}\left[\bar{\epsilon}(k)\right] - K(k)z_{\mathrm{m}}(kT_s) \\
P(k) &\triangleq \mathbb{E}\left[\hat{\epsilon}(k)\hat{\epsilon}^T(k)\right] = (I - K(k)H)\bar{P}(k),
\end{aligned}$$

where

$$\begin{aligned}
F &= e^{AT_s}, \\
Q &= \int_0^{T_s} e^{A\tau}\Gamma Q_c \Gamma^{\mathrm{T}} e^{A^{\mathrm{T}}\tau} d\tau, \\
S(k) &= H\bar{P}(k)H^{\mathrm{T}} + R, \\
K(k) &= \bar{P}(k)H^{\mathrm{T}}S^{-1}(k), \text{ and} \\
\mathbb{E}\left[\bar{\epsilon}(0)\right] &= 0.
\end{aligned}$$

Moving forward, it is assumed that the estimation error covariances have reached their steady-state values, which can be found by solving a discrete-time algebraic Riccati equation, and so the index k is dropped from P, $\bar{P}$, S, and K. Note that during a spoofing attack, a nonzero $z_{\mathrm{m}}(kT_s)$ biases the estimation error and innovation.

The NIS detection statistic $q(k) \triangleq \nu^{\mathrm{T}}(k)S^{-1}\nu(k)$ is distributed under H_0 as χ^2 with two degrees of freedom, as shown by [39], given $\nu(k)$ is a 2×1 vector. Similarly, under H_1, $q(k)$ is distributed as non-central χ^2 with two degrees of freedom and noncentrality parameter $\delta(k) = \bar{\nu}^{\mathrm{T}}(k)S^{-1}\bar{\nu}(k)$, where

$$\bar{\nu}(k) \triangleq H\mathbb{E}\left[\bar{\epsilon}(k)\right] - z_{\mathrm{m}}(kT_s)$$

is the mean of innovation at index k. Since the innovation sequence is white, each detection test is independent when conditioned on H_0 or H_1, which simplifies calculation of $\bar{I}_R$ because the integrand in (3) can be written as the product of the marginal probabilities for the events $q(k) < \lambda$, $k \in S_L$.

Optimization

A natural question arises in sequential detection: How often should the detection test be executed? In general, the detection test could be based on M measurements over a detection interval T_{det}, i.e., $T_{\mathrm{det}} = MT_s$. An initial simulation study has shown that, for the problem studied here, $M = 1$ yields the most powerful test for a given mean time between false alarms M_F. Thus, T_s will be taken as equivalent to T_{det} hereafter. An analytical proof that $M = 1$ is optimum remains an open problem. If T_s, is too small, then no innovation $\nu(k)$ will appear particularly surprising under H_1. As T_s is made longer, innovations under H_1 become more obviously biased. But if T_s is too long, the attacker may begin an attack and achieve his goal of reaching hazardous conditions all within the span between consecutive detection tests.

A distinguishing feature of the current framework is that it optimizes T_s to minimize $\bar{I}_R$ over a range of possible $v_{\max}$. Figure 5 shows how $\bar{I}_R$ varies as a function of T_s and $v_{\max}$ for an example scenario. The optimal T_s is a minimax solution which minimizes the maximum $\bar{I}_R$ over the range of $v_{\max}$ considered, in this case 0.1 m/s to 1 m/s. More formally, a robust

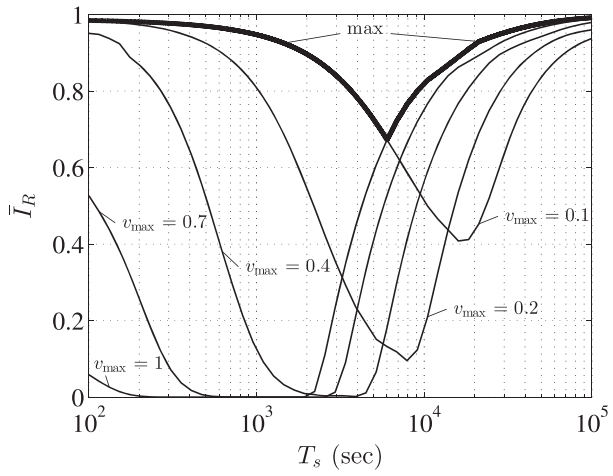

Fig. 5–Mean integrity risk $\bar{I}_R$ vs. sampling time T_s for various choices of $v_{\max}$. The optimal sampling time $T_s^\star$ that minimizes the worst-case mean integrity risk is approximately 100 minutes, yielding $\bar{I}_R^\star \approx 0.6727$. Note that the worst-case attack is given by either $v_{\max} = 0.1$ or 1 m/s. Other parameters are $u_{\max} = 0.03$ m/s^2, $M_F = 1$ month, $\bar{e} \gg L = 3$ km, $\sigma_p = 6$ m, $T_d = 200$ s, and $\sigma_d = 0.02$ m/s$^{1.5}$.

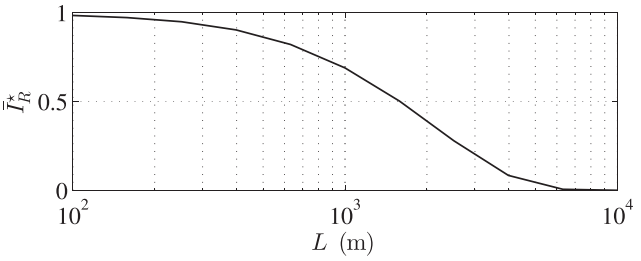

Fig. 6–Minimax mean integrity risk $\bar{I}_R^\star$ vs. the hazardous condition threshold L. For $L \leq 400$ m, the worst-case attack will likely cause HMI since $\bar{I}_R^\star > 0.9$. On the other hand, $L \geq 7$ km maintains an integrity risk near zero for any reasonable attack. Other parameters are set to the values indicated in Figure 5.

optimizer for T_s would be

$$\min_{T_s} \max_{\substack{v_{\max} \in \mathbb{V} \\ u_{\max} \in \mathbb{U}}} \bar{I}_R, \tag{4}$$

where $\mathbb{V}$ and $\mathbb{U}$ are bounded sets containing reasonable values for the attack parameters. $v_{\max}$ is bounded from below under the assumption that the attacker wishes to cause hazardous conditions before the end of a typical approach. If the average duration of an approach is $\bar{T}_{\text{app}}$, then $v_{\max} \geq L/\bar{T}_{\text{app}}$, reasonably assuming a linear relationship between the approach's alert limit and average duration. $v_{\max}$ is bounded from above because induced velocities greater than 1 m/s would lead to physically impossible set and drift values that are not captured by the Gauss-Markov disturbance model and break the small control error assumption. Lastly, the impact of $u_{\max}$ on $\bar{I}_R$ is small for the T_s values considered because acceleration only affects the attack profile for a short period of time in the beginning of the attack. Therefore, the integrity risk optimization is not particularly sensitive to the choice of $u_{\max}$, which is fixed to a value of 0.03 m/s^2 for the rest of the analysis.

A closed form solution to (4) does not appear possible, but the optimal T_s can be found numerically based on the definition of $\bar{I}_R$ and on the known distributions for $q(k)$ under H_0 and H_1. Minimax results for two example scenarios are shown in Figures 6 and 7.

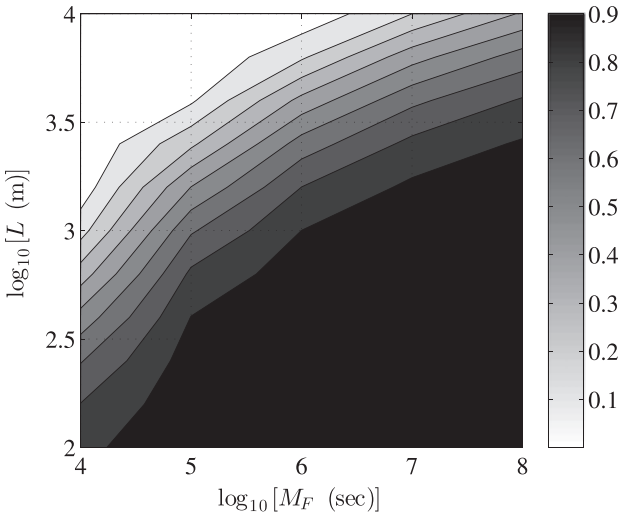

Fig. 7–Minimax mean integrity risk $\bar{I}_R^\star$ vs. L and M_F. Depending on the alert limit and continuity risk requirements of the approach, the detector's mean integrity risk ranges from high (black region), in which HMI is likely under H_1, to low (white region). Other parameters are set to the values indicated in Figure 5.

SIMULATION

The spoofer control law and integrity risk calculations were verified with Monte-Carlo simulations. The simulations take into account the Nomoto ship model, closed-loop ship controller, and open-loop spoofer controller developed in Section III, with $\bar{e} \gg L$. Two representative ship trajectories are shown in Figure 8. The simulation-based mean integrity risk is determined by counting the number of HMI events over 20 simulations with random measurement and process noise per attack profile and over 100 uniformly-spaced sampling phases per simulation. As shown in Figure 9, the simulation-based mean integrity risk for different values of $v_{\max}$ agrees well with the values predicted by the theory developed in Section IV.

DEMONSTRATION

Hostile control of a surface vessel by GPS spoofing was demonstrated in the Mediterranean Sea in June of 2013. The authors were invited to conduct

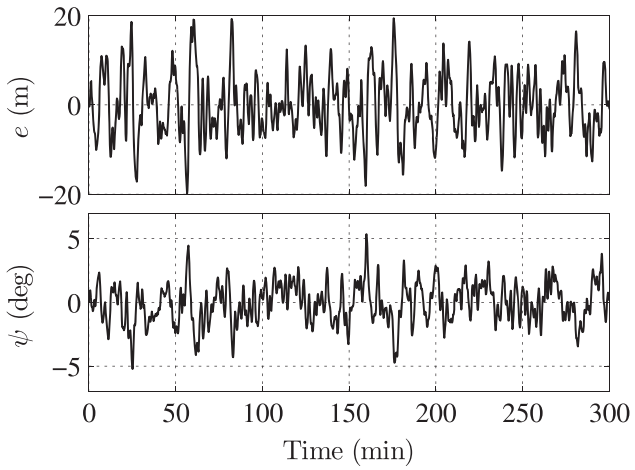

(a) Case I, no spoofing

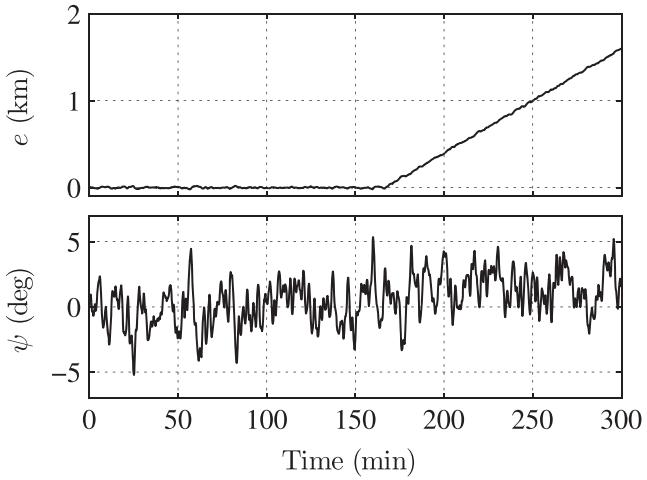

(b) Case II, $v_{\max} = 0.2$ m/s

Fig. 8–Trajectory resulting from simulation of ship dynamics under nominal conditions (a) and a spoofing attack (b). e and ψ are the ship's cross-track position and heading, respectively. Note that under the spoofing attack, there is a slight change in the average heading after the attack begins. The time-correlated heading offset may not look unusual to the crew depending on the expected time constants of ocean currents or wind in the area. Model parameters were $T = 39.94$ s, $K = 0.211\,\mathrm{s}^{-1}$, $U = 8.23$ m/s, $K_\mathrm{p} = 1.4415$, $K_\mathrm{i} = 0.0126$, $K_\mathrm{d} = 21.6904$, $K'_\mathrm{p} = 0.0028$, $K'_\mathrm{i} = 1.8949 \times 10^{-5}$. *Other parameters were set to the values indicated in Figure 5.*

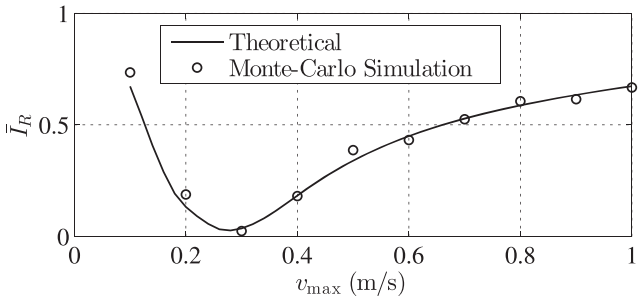

Fig. 9–Theoretical vs. simulated mean integrity risk for different values of $v_{\max}$. *Other parameters were set to the values indicated in Figure 5.*

the unprecedented experiment aboard the White Rose of Drachs, a 65-m superyacht. A key component of the experimental setup was a portable GPS spoofing device developed at the University of Texas at Austin [14]. The spoofer continuously received authentic GPS signals from an antenna on the ship's upper aft deck, and transmitted counterfeit GPS signals toward the ship's GPS antennas, located above the bridge as shown in Figure 10.

Once a safe route was established, the captain and his deck crew piloted the ship within a prescribed corridor along a series of rhumb lines. Periodic control actions were required to maintain course due to such disturbances as wind and ocean current, which were not measured directly. Instead, a lumped set and drift were measured indirectly from gyrocompass, Doppler speed log, and GPS measurements.

The spoofing attack was designed to cause a cross-track drift in the ship's apparent position that could be explained as the effect of an ocean current. The experiment was composed of three stages: (1) a subtle-attack stage with spoofer-induced cross-track velocity $v = 0.5$ m/s and $u_{\max} = 0.03$ m/s^2; (2) an aggressive-attack stage starting at $e = 200$m with $v = 2$ m/s and $u_{\max} = 0.1$ m/s^2; and (3) a parallel-track stage starting at $e = 700$ m during which v was reset to zero with $u_{\max} = 0.1$ m/s^2. Throughout the attack, the captain performed correction maneuvers to maintain the apparent (spoofed) ship position

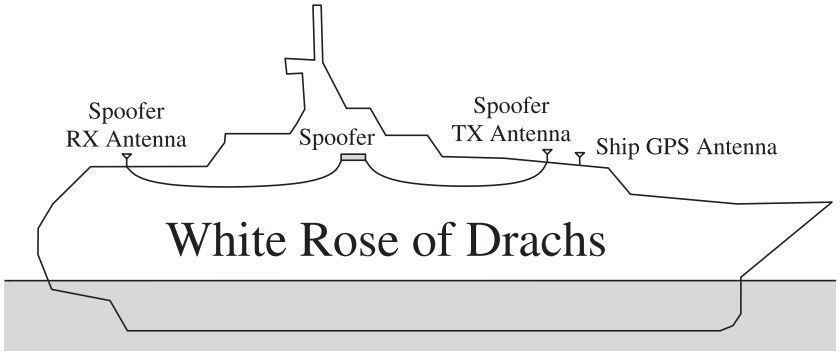

Fig. 10–Sketch of the spoofer setup on the White Rose of Drachs.

within a ±200 m corridor; the ship's actual position diverged along the spoofer-intended track shown in Figure 11.

The ship's GPS-reported position and gyrocompass - reported heading were logged to a file during the spoofing attack. The ship's Doppler log was not functional during the experiment, but the ship's engine throttle control was held constant at "full ahead," so the ship's speed through water U was assumed to be a nominal 15 knots. The ship's true position was computed by the spoofer using the authentic signals received by an antenna with sufficient isolation from the spoofed signals as shown in Figure 10.

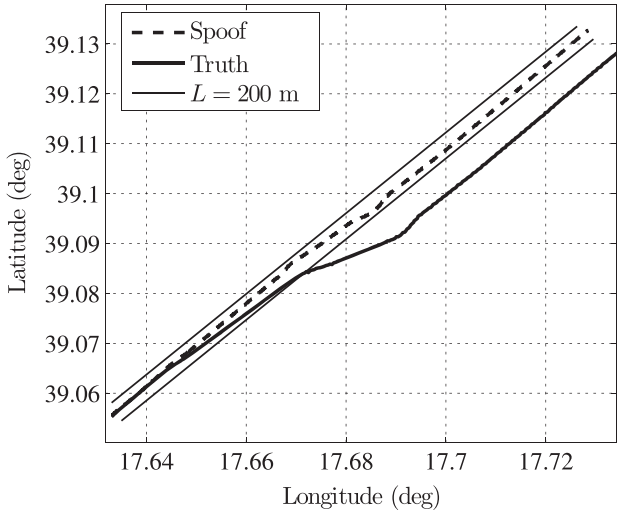

Fig. 11–Comparison of the ship's reported actual position during the spoofing attack. The thin solid lines mark a ±200 m safe corridor.

The logged measurements were fed post-facto into the innovations-based spoofing detector developed in Section IV. To determine the optimal sampling time $T_s^\star$ for the experiment, many of the same parameter values indicated in Figure 5 were used, except that 0.5 m/s $\leq v_{\max} \leq$ 2 m/s and L = 200 m. Even though the ship was traveling in open waters, a narrow "safe" corridor was chosen to simulate a situation with tight maneuverability bounds such as a harbor approach bordered by underwater hazards. The resulting minimax optimization yielded $T_s^\star \approx$ 250 s and mean integrity risk $\bar{I}_R^\star$ = 0.8956 for the worst-case attacks. The subtle attack demonstration shows that for an alert limit less than 200 m, the navigation solution from the sampling-time-optimized system described in the paper becomes unavailable for any mean integrity risk specification less than 0.89. If the system is unavailable, the navigator can improve the integrity risk of the existing system by trying to improve the model (i.e., reduce uncertainty), increase the continuity risk limit, or use another system to navigate the approach.

The attack profile applied during the subtle-attack stage was designed to be a worst case; unsurprisingly, the attack remained undetectable during this stage. The aggressive-attack stage was designed to be obvious: even assuming the inflated parameters L = 700 m, $u_{\max}$ = 0.1 m/s^2, and $v_{\max}$ = 2 m/s, the theoretical mean integrity remained an insignificant

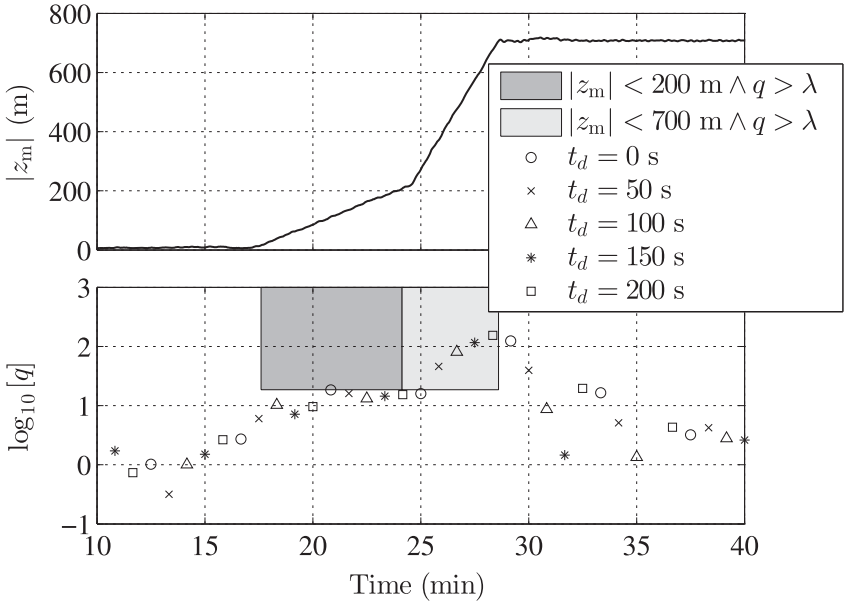

Fig. 12–NIS values generated by the detector with sampling time T_s = 250 s for the experimental data collected on the White Rose of Drachs during a live spoofing attack. A particular sampling phase t_d represents the delay since the experiment start time for the first detection test. NIS time histories for five different sampling phases are shown. The shaded regions mark areas where the NIS must fall for the attack to be detected before hazardous conditions occur, preventing an HMI event. The darker and lighter regions correspond to the first and second stages of the attack, respectively. The lower edge of the regions corresponds to the detection threshold λ.

$\bar{I}_R$ = 0.0067. The actual integrity risk was different due to the changes in the spoofer-induced velocity during the attack.

The NIS values generated by the detector based on experimental data with five different sampling phases are shown in Figure 12. Recall that the mean integrity risk computed previously is the marginal risk assuming a uniformly-distributed sampling phase. A realization for a particular sampling phase leads to HMI if the associated NIS values fail to cross the detection threshold λ before an attack reaches hazardous conditions. The NIS values in Figure 12 fail to cross into the dark shaded region before $z_{\mathrm{m}} > 200$ m (equivalently, $e > 200$ m), indicating that the attack was not detected before hazardous conditions during the subtle stage. The sampling phase t_d can be related to start time of the spoofing attack as $t_d \triangleq \tau_1 - t_0$, where τ_1 is the time of the first detection test after the onset of spoofing. For all but one sampling phase ($t_d = 0$ s), the attack was detected during the aggressive stage before hazardous conditions occurred.

SENSOR-LEVEL STRATEGIES FOR MITIGATING SURFACE VESSEL VULNERABILITY TO GNSS DECEPTION

Despite being tailored to minimize $\bar{I}_R$, the detector developed in this paper remains vulnerable to subtle spoofing attacks that masquerade as the effect of ocean currents, as was demonstrated vividly by the experiment aboard the White Rose of Drachs. What is more, it is doubtful that any other detector operating on measurements from a GNSS receiver and from the standard dead-reckoning instruments (the gyrocompass and Doppler speed log)—or any other navigation sensors common to today's surface vessels—could detect a subtle GNSS spoofing attack before hazardous conditions occur.

The difficulty of reliable spoofing detection at the sensor fusion level motivates a layered approach in which the detector proposed in this paper is complemented with a GNSS receiver also designed to detect spoofing. A number of promising receiver-level spoofing detection methods are surveyed in [40]. Among these, the dual-antenna technique advanced in [41] seems an especially promising option for maritime protection because (1) it can be implemented in the near term, and (2) its chief drawbacks relative to the other techniques—larger size and higher cost—are not so critical for marine vessels as they are for handheld devices and small unmanned aerial vehicles, for example. Nonetheless, it will take years before this or other techniques mature and are implemented widely. Meanwhile, there are no off-the-shelf defenses against GNSS spoofing.

CONCLUSIONS

Modern integrated bridge systems assume a surface vessel's GNSS receivers are trustworthy when these report a position fix from ambient GNSS signals. But such trust is misplaced in situations of GNSS spoofing: spoofed receivers report an attacker-induced false ship position as conveyed via counterfeit GNSS signals. An attacker can modulate the ship's true along-track and cross-track positions by feeding apparent positions to the ship's autopilot system, or to its bridge crew, that are falsely offset from the ship's true position. Besides this system-level effect of spoofing, specific navigation and collision avoidance instruments are individually affected: the automatic radar plotting aid, the automatic identification system, the dead reckoning system built into the ship's electronic chart display and information system (ECDIS), and the ship's satellite compass can all generate hazardously misleading information during a GNSS spoofing attack.

A detection framework was developed to analyze and detect spoofing attacks on surface vessels based solely on Doppler log, gyrocompass, and GNSS measurements. The framework's detector is implementable in ECDIS software commonly available on ships of significant size. The detector is based on a dynamics model that captures the essential features of the environmental disturbances, which are dominated by ocean currents and wind. The detector's test interval was chosen to minimize the maximum mean integrity risk, or probability of hazardously misleading information averaged over possible attack start times, for a range of goal-oriented attack profiles. Monte-Carlo simulations verified the theoretical calculations of mean integrity risk. An unprecedented experiment demonstrated successful hostile control of an actual surface vessel in a live spoofing attack and detection of the attack during its most aggressive stage.

Just as aviation regulators have developed rigorous integrity risk standards for GNSS faults, maritime regulatory authorities can use the detection framework analysis proposed herein to compute the minimum integrity risk given reasonable values for real-world disturbance and attack parameters and the maximum acceptable continuity risk.

ACKNOWLEDGEMENTS

This work was supported in part by the National Science Foundation under Grant No. 1454474 and by the Data-supported Transportation Operations and Planning Center (D-STOP), a Tier 1 USDOT University Transportation Center. The authors thank the crew of the White Rose of Drachs and her captain, Andrew Schofield, for supporting the

GPS spoofing experiment. Thanks also to Dr. Andy Norris for his comments on maritime navigation and security in practice.

REFERENCES

1. John A. Volpe National Transportation Systems Center, "Vulnerability Assessment of the Transportation Infrastructure Relying on the Global Positioning System," 2001.
2. Grant, A., "GPS Jamming and the Impact on Maritime Navigation," *Journal of Navigation*, Vol. 62, No. 2, 2009.
3. International Marine Contractors Association, "Guidelines for the Design and Operation of Dynamically Positioned Vessels," 2007. Available at: http://www.imca-int.com/media/73055/imcam103.pdf.
4. Thomas, M., Norton, J., Jones, A., Hopper, A., Ward, N., Cannon, P., Ackroyd, N., Cruddace, P., and Unwin, M., "Global Navigation Space Systems: Reliance and Vulnerabilities," *The Royal Academy of Engineering*, London, 2011.
5. Caccia, M., Bibuli, M., Bono, R., and Bruzzone, G., "Basic Navigation, Guidance and Control of an Unmanned Surface Vehicle," *Autonomous Robots*, Vol. 25, No. 4, 2008, pp. 349–365.
6. Elkins, L., Sellers, D., and Monach, W. R., "The Autonomous Maritime Navigation (AMN) Project: Field Tests, Autonomous and Cooperative Behaviors, Data Fusion, Sensors, and Vehicles," *Journal of Field Robotics*, Vol. 27, No. 6, 2010, pp. 790–818.
7. Paull, L., Saeedi, S., Seto, M., and Li, H., "AUV Navigation and Localization: A Review," *IEEE Journal of Oceanic Engineering*, Vol. 39, No. 1, 2014, pp. 131–149.
8. Meduna, D. K., Rock, S. M., and McEwen, R. S., "Closed-Loop Terrain Relative Navigation for AUVs with Non-Inertial Grade Navigation Sensors," *Proceedings of 2010 IEEE/OES Autonomous Underwater Vehicles (AUV)*, 2010, pp. 1–8.
9. Meduna, D., Rock, S. M., and McEwen, R., "AUV Terrain Relative Navigation using Coarse Maps," *Proceedings of Unmanned Untethered Submersible Technology Conference*, 2009.
10. U.S. Coast Guard; U.S. Department of Homeland Security, "Terminate Long Range Aids to Navigation (Loran-C) Signal," *Federal Register*, January 2010.
11. Narins, M., Lombardi, M., Enge, P., Peterson, B., Lo, S., Chen, Y. H., and Akos, D., "The Need for a Robust Precise Time and Frequency Alternative to Global Navigation Satellite Systems," *Journal of Air Traffic Control*, Vol. 55, No. 1, 2012.
12. Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., O'Hanlon, B. W., and Kintner, P. M. Jr., "Assessing the Spoofing Threat: Development of a portable GPS Civilian Spoofer," *Proceedings of the 21st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2008)*, Savannah, GA, September 2008, pp. 2314–2325.
13. Shepard, D. P., Humphreys, T. E., and Fansler, A. A., "Evaluation of the Vulnerability of Phasor Measurement Units to GPS Spoofing Attacks," *International Journal of Critical Infrastructure Protection*, Vol. 5, No. 3–4, 2012, pp. 146–153.
14. Kerns, A. J., Shepard, D. P., Bhatti, J. A., and Humphreys, T. E., "Unmanned Aircraft Capture and Control via GPS Spoofing," *Journal of Field Robotics*, Vol. 31, No. 4, 2014, pp. 617–636.
15. GPS Directorate, "Systems Sngineering and Integration Interface Specification IS-GPS-200G," 2012. Available at: http://www.gps.gov/technical/icwg/.
16. European Union, "European GNSS (Galileo) Open Service Signal in Space Interface Control Document," 2010. Available at: http://ec.europa.eu/enterprise/policies/satnav/galileo/open-service/.
17. Willsky, A. S., "A Survey of Design Methods for Failure Detection in Dynamic Systems," *Automatica*, Vol. 12, No. 6, 1976, pp. 601–611.
18. Basseville, M., "Detecting Changes in Signals and Systems—a Survey," *Automatica*, Vol. 24, No. 3, 1988, pp. 309–326.
19. Frank, P. M., "Fault Diagnosis in Dynamic Systems Using Analytical and Knowledge-Based Redundancy: A Survey and Some New Results," *Automatica*, Vol. 26, No. 3, 1990, pp. 459–474.
20. Chen, J., and Patton, R. J., *Robust Model-Based Fault Diagnosis for Dynamic Systems,* Springer Publishing Company, Inc., 2012.
21. Joerger, M., and Pervan, B., "Kalman Filter-Based Integrity Monitoring Against Sensor Faults," *Journal of Guidance, Control, and Dynamics*, Vol. 36, 2013, pp. 349–361.
22. Khanafseh, S., Roshan, N., Langel, S., Cheng-Chan, F., Joerger, M., and Pervan, B., "GPS Spoofing Detection using RAIM with INS Coupling," *Proceedings of the IEEE/ION PLANS 2014*, Monterey, CA, May 2014, pp. 1232–1239.
23. Wald, A., "Sequential Tests of Statistical Hypotheses," *The Annals of Mathematical Statistics*, Vol. 16, No. 2, June 1945, pp. 117–186.
24. Trees, H. L. V., *Detection, Estimation, and Modulation Theory,* Wiley, 2001.
25. Mehra, R. K., and Peschon, J., "An Innovations Approach to Fault Detection and Diagnosis in Dynamic Systems," *Automatica*, Vol. 7, No. 5, 1971, pp. 637–640.
26. Pelkowitz, L., and Schwarts, S., "Asymptotically Optimum Sample Size for Quickest Detection," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-23, No. 2, March 1987, pp. 263–272.
27. Bowditch, N., *The American Practical Navigator*, Bethesda, MD, National Imagery and Mapping Agency, 2002.
28. GPS World Staff, "Hemisphere GPS Offers Vector Compass Products for Marine Applications," Oct. 2012. *GPS World*, Available at: http://gpsworld.com/hemisphere-gps-offers-vector-compass-products-for-marine-applications.
29. *Radar Navigation and Maneuvering Board Manual*, 7th ed., Bethesda, Maryland: National Imagery and Mapping Agency, 2001. Available at: http://msi.nga.mil/MSISiteContent/StaticFiles/NAV_PUBS/RNM/310ch5.pdf.
30. Norris, A., *ECDIS and Positioning*, Series, Integrated Bridge Systems: Nautical Institute, 2010.
31. eNav International, "Totem ECDIS and GPS Spoofing," June 2013. Available at: http://www.enav-international.com/news/id5774-Totem_ECDIS_and_GPS_Spoofing.html.
32. Fossen, T. I., *Guidance and Control of Ocean Vehicles*, New York: John Wiley and Sons, 1994.
33. National Transportation Safety Board, "Marine Accident Report: Grounding of the Panamanian Passenger Ship Royal Majesty on Rose and Crown Shoal near Nantucket, Massachusetts June 10, 1995," *Technichal Report,* National Transportation Safety Board, 1997.

34. Lützhöft, M. H., and Dekker, S. W., "On Your Watch: Automation on the Bridge," *Journal of Navigation*, Vol. 55, No. 1, 2002, pp. 83–96.
35. Kendoul, F., "Survey of Advances in Guidance, Navigation, and Control of Unmanned Rotorcraft Systems," *Journal of Field Robotics*, Vol. 29, No. 2, 2012, pp. 315–378.
36. Bhatti, J., "Sensor Deception Detection and Radio-Frequency Emitter Localization," Ph.D. Dissertation, The University of Texas at Austin, August 2015.
37. Luce, R. D., and Raiffa, H., *Games and Decisions: Introduction and Critical Survey,* Dover, 1989.
38. Blackwell, D. A., *Theory of Games and Statistical Decisions,* Courier Dover Publications, 1979.
39. Bar-Shalom, Y., Li, X. R., and Kirubarajan, T., *Estimation with Applications to Tracking and Navigation*, New York: John Wiley and Sons, 2001.
40. Humphreys, T. E., *The GNSS Handbook,* Springer, 2014. ch. Interference, in preparation.
41. Psiaki, M. L., O'Hanlon, B. W., Powell, S. P., Bhatti, J. A., Wesson, K. D., Humphreys, T. E., and Schofield, A., "GNSS Spoofing Detection using Two-Antenna Differential Carrier Phase," *Proceedings of the* 27^{th} *International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2014)*, Tampa, FL, September 2014, pp. 2776–2800.