

This work was written as part of one of the author's official duties as an Employee of the United States Government and is therefore a work of the United States Government. In accordance with 17 U.S.C. 105, no copyright protection is available for such works under U.S. Law.

Public Domain Mark 1.0


<https://creativecommons.org/publicdomain/mark/1.0/>

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Transitioning from testbeds to ships: an experience study in deploying the TIPPERS Internet of Things platform to the US Navy

Journal of Defense Modeling and Simulation: Applications, Methodology, Technology
2022, Vol. 19(3) 501–517
© The Author(s) 2020
DOI: 10.1177/1548512920956383
journals.sagepub.com/home/dms


Dave Archer⁵, Michael A August³, Georgios Bouloukakis¹, Christopher Davison^{1,2} , Mamadou H Diallo³, Dhrubajyoti Ghosh¹, Christopher T Graves³, Michael Hay⁸, Xi He⁷, Peeter Laud⁶, Steve Lu⁴, Ashwin Machanavajjhala¹⁰, Sharad Mehrotra¹, Gerome Miklau⁹, Alisa Pankova⁶, Shantanu Sharma¹, Nalini Venkatasubramanian¹, Guoxi Wang¹ and Roberto Yus¹

Abstract

This paper describes the collaborative effort between privacy and security researchers at nine different institutions along with researchers at the Naval Information Warfare Center to deploy, test, and demonstrate privacy-preserving technologies in creating sensor-based awareness using the Internet of Things (IoT) aboard naval vessels in the context of the US Navy's Trident Warrior 2019 exercise. Funded by DARPA through the Brandeis program, the team built an integrated IoT data management middleware, entitled TIPPERS, that supports privacy by design and integrates a variety of Privacy Enhancing Technologies (PETs), including differential privacy, computation on encrypted data, and fine-grained policies. We describe the architecture of TIPPERS and its use in creating a *smart ship* that offers IoT-enabled services such as occupancy analysis, fall detection, detection of unauthorized access to spaces, and other situational awareness scenarios. We describe the privacy implications of creating IoT spaces that collect data that might include individuals' data (e.g., location) and analyze the tradeoff between privacy and utility of the supported PETs in this context.

Keywords

Internet of Things, privacy, US Navy, deployment, practical experiences, secure computing, differential privacy, data management, middleware

1. Introduction

Advances in sensing, networking, and communication technologies have created a new wave of Internet of Things (IoT) technologies that are expected to revolutionize all aspects of our society. In the context of armed forces, the ability to dynamically monitor soldiers, their physiological health possibly using biometric wearables, and the environment in which they operate can be used to build a variety of military applications. IoT technologies can potentially connect, in real time, ships, planes, tanks, drones, personnel (both on board a ship or on land) for creating a cohesive fighting force with improved situational awareness. The overall objective is to bring about

¹University of California, Irvine, USA

²Ball State University, USA

³Naval Information Warfare Center, Pacific, USA

⁴Stealth Software Technologies, Inc., USA

⁵Galois Inc., USA

⁶Cybernetica, Estonia

⁷University of Waterloo, Canada

⁸Colgate University, USA

⁹University of Massachusetts, Amherst, USA

¹⁰Duke University, USA

Corresponding author:

Christopher Davison, Ball State University, Muncie, IN 47306-1022, USA.
Email: cbdavison@bsu.edu

transformative improvements in the ability of the soldiers to improve battlefield performance and outcomes. IoT technologies can play a significant role in enhancing the ability of soldiers not just by creating real-time awareness; multimodal sensor data collected during missions from diverse sources can also be analyzed post-facto to extract and evaluate the operational aspects of mission goals leading to future process improvements.

While the importance of the emerging IoT technologies and applications cannot be overemphasized, a key challenge facing their widespread adoption is that of privacy. The continuously captured sensor data, which can leak information about subjects, their habits, likes/dislikes, health, mental status, etc., has been well studied. For instance, our early deployment studies at a building at the University of California, Irvine (UCI), clearly established that even coarse-level monitoring of subjects using their connectivity to WiFi Access Points (APs) in a building can lead to personal information leakage, such as the amount of time people are at work, the number of times they take a break during the day, tardiness/effectiveness in performing their tasks, etc. Another study, at the Honeywell Golden Valley Labs,¹ further corroborated such privacy-sensitive information leakage; here, monitoring seemingly innocuous sensor data from motion sensors (for energy-efficient building operation) coupled with background knowledge could lead to determining the personal habits of individuals, including smoking habits and arrival/departure times from work.

The layman's perspective on privacy is that it is largely a civilian concern, since individuals who enlist in the military relinquish privacy rights in the interest of national security. Our experience in the context of this study demonstrates that this is not true in day-to-day operational circumstances. Privacy is critical at all levels of military command and control—from senior leadership down the chain of command to enlisted sailors/soldiers. From an executive point of view, concerns arise about eroding the morale of personnel who perceive that they are constantly monitored (big brother syndrome)—this raises issues of long-term retention of personnel. At lower ranks of the military, there are concerns that privacy is a scarce resource when individuals are required to stay in close proximity to each other for elongated periods of time. Any personal time that is afforded to them during their off-duty hours is a cherished resource; monitoring personal activity and behavior through sensors continuously without a compelling or immediate reason during those times can be a concern in terms of technology adoption.

There is yet another important rationale for privacy from a tactical perspective, as studied in prior work.² In his research on radio frequency identification (RFID) signaling specifically, Juels² (p.5) states “privacy is not just a consumer concern,” and he elaborates on the topic,

providing examples of enemy forces harvesting radio frequency (RF) signals to ascertain troop movements and logistical support activities. Furthermore, Jules states that RFID tags detecting munitions are a plausible threat. From this perspective in particular, it becomes apparent that privacy often equates to security, safety, and victory.

It is also worth noting that troops often purchase personal items from the same retail outlets as civilians. Along with that, such goods routinely have RFID or other RF assets embedded within (e.g., Bluetooth or WiFi). According to Nayak,³ many retailers, including Walmart, mandate their suppliers to use RFID technology. Therefore, it is quite possible that rogue or unwanted RFID equipment can potentially maneuver its way into the battlespace.

The battlespace and its concomitant warriors are becoming increasingly saturated with technology. Provisioning and utilization of new technologies is necessary in order to enhance situational awareness, support force effectiveness, increase military efficiency, and maintain competitive advantage in a technology-augmented global environment.

Considering this need for data privacy, DARPA initiated four years ago the Brandeis program,⁴ which seeks to develop the technical means to protect the private and proprietary information of individuals and enterprises. The main goal of the program is to develop tools and techniques that enable systems to be built with privacy in mind. The *Testbed for IoT-based Privacy-preserving Pervasive Spaces* (TIPPERS) system is part of this program. TIPPERS addresses the program's goal by exploring a new generation of emerging privacy technologies, including encryption, multi-party computation (MPC), differential privacy (DP), and privacy policies, as a basis for building an IoT data collection and management system that supports *privacy by design*. TIPPERS was initially designed and tested at the University of California Irvine, where it has been used to create a smart campus environment with a variety of applications, such as space occupancy monitoring, ability to locate friends, etc.

Over the past year, this research team has worked toward transforming TIPPERS to be deployed in tactical naval settings. In particular, TIPPERS was deployed on a Navy ship as part of the annual Trident Warrior exercise in 2019. According to the US Navy, “Trident Warrior is an annual large-scale, at-sea field experiment where the Navy selects potential initiatives that address capability gaps and provide inventive solutions in an operational environment” (Military News website,⁵ (para.2)).

This paper highlights the team's research experience in deploying TIPPERS during the Trident Warrior 2019 (TW 19) exercise. The potential benefits of the IoT technologies for naval use cases, the concerns related to privacy that emerged, and possible approaches to address those

concerns are presented. Our results indicate that warfighters that are outfitted with these often-invasive technologies tend to experience a lack of privacy. Such infringement can result in adverse effects on morale, quality of life, and personnel retention. However, these same technologies bring efficiencies to command and control efforts, enhance situational awareness, and make warfighters more safe and secure.

The rest of the paper is organized as follows. Section 2 provides a brief overview of the TIPPERS system and underlying privacy technologies. Section 3 discusses the TIPPERS deployment on the ship, modeling, and simulation of sailor activity, as well as the instrumentation of both the ship and the individuals. Section 4 outlines specific aspects of the use cases/scenarios designed for the Trident Warrior exercise and experiments, focusing on privacy studies in Section 5. Finally, a discussion of lessons learned and experiences working with the US Navy concludes the paper in Section 6.

2. TIPPERS

TIPPERS is a novel sensor data collection and management system for smart spaces that incorporates a variety of smart space applications.⁶⁻⁸ A key design feature of the TIPPERS architecture is that it is space, sensor, and task

agnostic, allowing it to be used as plug-and-play technology to create smart spaces. In addition, TIPPERS embodies a *privacy-by-design* architecture, which enables the integration of different Privacy Enhancing Technologies (PETs). In particular, and in the context of the DARPA Brandeis Project, a variety of PETs have been integrated, including secure computing, privacy policies, and DP (see Appendix 1 (supplementary material) for more information about such PETs).

2.1 TIPPERS design

As depicted in Figure 1, the TIPPERS architecture includes several decisions to support the goal of privacy by design. Firstly, TIPPERS provides an abstraction of the underlying sensor infrastructure by translating between the *IoT devices' world* (i.e., sensors, actuators, raw observations, etc.) and the *people's world* (i.e., interactions of people, spaces, phenomena, etc.). The system is based on a domain model that represents both worlds and enables users/developers to interact with high-level semantically meaningful concepts. It also includes ontology-based translation algorithms to convert user requests at the high level (e.g., “decrease the temperature of rooms where the occupancy is greater than 75% of their capacity”) into actions on the specific underlying device infrastructure.⁷ The main

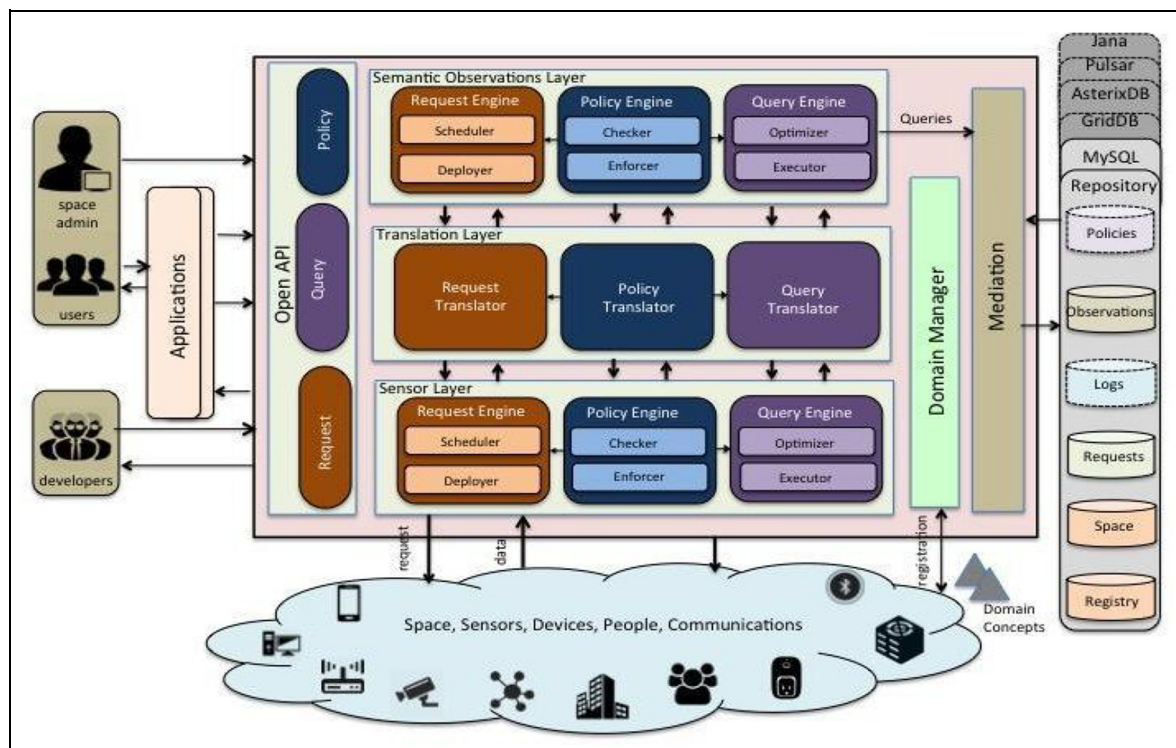


Figure 1. Architecture of TIPPERS. API: application programming interface.

advantage is that it simplifies the development of smart applications and facilitates their portability in between spaces as they are built on high-level concepts instead of on IoT devices. Secondly, it simplifies the definition of privacy policies as users can focus on what they want to protect (e.g., “do not capture my location when I am with John in a private space during working hours”). TIPPERS uses such privacy policies to guide its data collection, storage, and sharing practices.⁸

As a mechanism to implement such translation of raw data into higher-level semantically meaningful interpretations, TIPPERS supports *virtual sensors* wherein streams of sensor data can be used to create streams of such inferences.⁹ For instance, a virtual sensor can translate connectivity data (e.g., logs from WiFi APs containing information about which devices are connected to them) into occupancy of different spaces along time. This enables TIPPERS to incorporate further PETs. For example, a stream of sensor data can be scrubbed of personally identifiable information (PII) when passed to operators.

Finally, the TIPPERS architecture contains a mediation module to appropriately store sensor data in the corresponding database/storage technology (e.g., allowing the usage of different underlying database systems). This mediation consists of three parts regarding the specific mapping task between the TIPPERS schema, data, and queries and those of the database system. This enables TIPPERS to store data in different systems based on the characteristics of the data, its security requirements, and the type of queries that need to be run on the data. Using this functionality, the TIPPERS system further includes secure data storage technologies that can maintain encrypted data and perform computations on such encrypted data.

2.2 Privacy Enhancing Technologies

Below we describe several PETs that were part of the TIPPERS system deployed in the US Navy ship (more information about the PETs is included in Appendix 1 (supplementary material)).

2.2.1 PULSAR. PULSAR is a novel secure data management system based on function secret sharing (FSS)^{10,11,18–22} and MPC^{12,13} to support real-time privacy-preserving data aggregation and retrieval that has been applied to sensors and mobile devices in TIPPERS. A standard secret-sharing scheme allows a dealer to randomly split a secret into two or more shares, such that certain subsets of the shares can be used to reconstruct the secret and others reveal nothing about it. Secret sharing is additively homomorphic, that is, if many secrets are shared, the two parties can individually compute shares of

the sum of the secrets by locally adding their shares, without any communication.

The notion of FSS can be viewed as a natural generalization of additive secret sharing to functions. A special case of interest is the class of point functions f , which have a nonzero output on at most one input. A FSS for this class is called a distributed point function (DPF). One exemplary application for such a special case is secure distributed histograms, or “distograms,” that allow for the ability to privately aggregate information into histogram buckets. Stealth (the makers of PULSAR) incorporates these distograms into the PULSAR solution as a means of real-time privacy-preserving data aggregation and retrieval.

2.2.2 Jana. Jana technology implements the paradigm of Private Data as a Service (PDaaS). Using a combination of advanced cryptographic techniques and a commercially reliable database, such technology provides a full-featured, robust, relational database management system (RDBMS). The RDBMS cryptographically secures data from before it leaves the platform of a data contributor, until after it reaches the platform of an analyst authorized to see the query results. Data does not need to be decrypted during query processing in Jana. Results of queries are additionally protected using DP mechanisms (where appropriate) to prevent rediscovery of sensitive data from those results.

Jana is also intended as an operational environment to study the trade-space between security and performance scalability for real-world data and queries. In contrast, typical cryptographic research platforms fix a level of security, argue on standard assumptions, and offer no trade-space in which to conduct such research. In addition, Jana allows for the study of implications on performance of full, end-to-end security, while typical cryptography research fails to address the security of all steps in the information flow from data provider to query result.

The data within Jana remains encrypted at all times, unless explicitly chosen by the database administrator (DBA) (and agreed to by data contributors) for storage in plaintext form. Ephemeral public key encryption protects data in transit from contributors’ platforms into each Jana instance, as well as results in transit from the Jana instance to analyst platforms. Depending on choices made by the DBA, public key encryptions or order revealing encryptions are used to protect data at rest in Jana’s relational data store (deterministic encryption is also supported as a research capability, although it is not recommended for the obvious security concerns).

2.2.3 PeGaSus. PeGaSus (PGS) is a specific algorithm for analyzing streaming data under DP.¹⁴ The technology of DP is appropriate in contexts where data about individuals has been curated for the purpose of analysis. The goal is to

release the outcome of the analysis while disclosing as little information as possible about individual records in the data. The conventional technologies for this problem fall under the broad category of *disclosure limitation* and include approaches such as data de-identification and suppression. Unfortunately, these techniques are known to be brittle. For instance, so-called “anonymous” records can often be re-identified. In contrast, DP offers a rigorous mathematical guarantee: to an individual whose data has been collected, whatever can be learned about the individual when her data is included is essentially no greater than what can be learned where her data is omitted from the collection.

DP is distinct from adjacent technologies, such as secure multi-party computation (SMC). With SMC, the goal is to compute the exact answer to a function where each party contributes one private input, with the goal that during the execution of the computation, no party is able to extract information about the inputs supplied by other parties. DP, on the other hand, aims to prevent the output of the function from leaking information. For instance, if the function computed a vote tally and it was unanimous, then the output of SMC would reveal individual votes; in contrast, with DP, one would learn only that the vote was approximately unanimous.

DP is a mathematical definition that can be applied to a variety of data types and for a variety of analyses. In considering its use in a particular application, it is essential to consider what is sensitive and private information, as well as what kinds of analyses should be supported.

PGS is applicable in contexts where individuals are continually observed by a collection of sensors. Each sensor produces a stream of data, and the goal is to analyze these streams in real time without disclosing sensitive personal information about specific individuals. With PGS, the privacy guarantee is on individual events, such as a person being observed by a particular sensor, but extends naturally to small windows of events (observations of a particular individual within the last hour).

2.2.4 Integration of PETs in TIPPERS. TIPPERS supports two ways of integrating PETs into sensor data management, as mentioned before. Firstly, virtual sensor technology supported by TIPPERS can be leveraged to modify/perturb the sensor data stream that is shared with applications. Such modifications could include scrubbing of sensitive data (e.g., removing faces from images or identifiers from sensor data), adding noise, de-linking data, etc. A good example of the use of virtual sensors is the integration of PGS to make occupancy data streams differentially private.²³

Secondly, TIPPERS supports mediation between the system and the underlying storage mechanism. For instance, TIPPERS mediates with FSS and MPC technologies supported by PULSAR, as well as deterministic

encryption, order-preserving encryption, and secret-sharing-based technologies supported by Jana. This way, TIPPERS can be configured to store low-level sensor data (such as WiFi connectivity) in PULSAR such that insertions can be fast and aggregations (e.g., occupancy levels) can be determined quickly. This is crucial for real-time policy enforcement when policies depend on who is in a particular space or how many people are inside it. In contrast, data that may require more complex operations (such as, for instance, Structured Query Language (SQL) joins operations to combine sensor data with metadata) can be stored in Jana, which supports different techniques for encryption and supports more complete complex SQL operations.

A sample flow of data in TIPPERS and the integrated PETs in the deployment in the Navy destroyer is illustrated in Figure 2, which is described in more detail in Appendix 1 (supplementary material).

2.3 TIPPERS deployment at University of California, Irvine

In addition to the deployment of TIPPERS aboard the Navy ship as part of the Naval Trident Warrior Exercises, the system has also been deployed at other locations. One location was the Golden Valley Lab at Honeywell, to create a privacy-preserving building analytics system using motion sensors. The primary testbed installation of TIPPERS is located at the UCI campus and is used to create a smart campus for UCI members. At UCI, TIPPERS is used daily at the campus-scale (over 30+ buildings) to support a variety of location-based services, such as detection of occupancy levels of buildings, building usage analysis, concierge services (allowing people to find each other on campus and inside buildings), and self-monitoring applications (allowing people to monitor how, where, and with whom they use their time). At the UCI deployment, TIPPERS determines the location of individuals using WiFi connectivity data from personal devices carried by individuals. Such connectivity data from WiFi APs is collected by the Office of Information Technology (OIT) in order to provide network services. In turn, such information is shared with TIPPERS after appropriately encrypting the MAC address of the devices to prevent leakage of a user’s location. The OIT changes the encryption key every 5 minutes in order to prevent linkage attacks. TIPPERS computes on top of such encrypted data to analyze, for instance, real-time occupancy at different granularities—viz., building, floors, regions within a floor—which is used to support a real-time heat map of the building.

3. TIPPERS in Trident Warrior 2019

In this section, we discuss the deployment of TIPPERS in the Navy ship as part of the US Navy’s TW 19 exercise.⁵

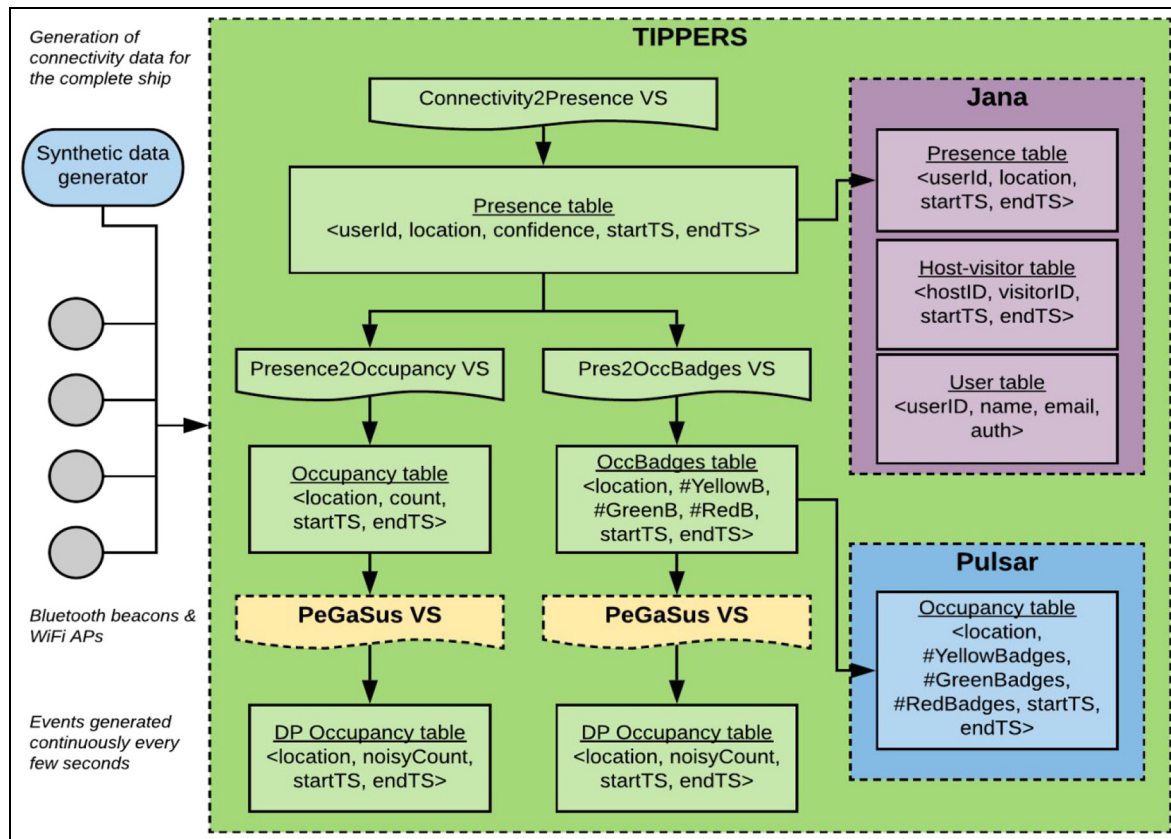


Figure 2. TIPPERS data flow. AP: Access Point; DP: differential privacy.

3.1 Testing TIPPERS for Trident Warrior deployment

As part of the preparation for TW 19, the TIPPERS system was deployed first at the Naval Information Warfare Center (NIWC) Pacific. The objective was to enable a demonstration of the capability of TIPPERS, including its integrated PETs, to operational Navy personnel to explore opportunities wherein systems such as TIPPERS can be used within the Navy. TIPPERS was deployed in the NIWC Mobility Center of Excellence (MCoE), which is a lab dedicated to developing, testing, and evaluating mobile technologies. The deployment required developing an application specific to NIWC, referred to as the *Security Surveillance* application.

3.1.1 Security Surveillance application. The *Security Surveillance* application provides a bird's eye view of the evolving state of the NIWC facility based on the sensor data that is captured and translated into occupancy levels of each space. The sensor data is securely encrypted in the underlying secure data storage TIPPERS technology. To enforce the *need to know* concept, the application provides two views of the data. The first shows only differentially private occupancy counts (obtained using the PGS virtual

sensor) with no identifying information to preserve the privacy of individuals involved in the data. From this view, a user of the application should not be able to make any inferences about individuals, their locations, or their habits. The second view shows the actual occupancy counts after decryption (as this data is stored in the PULSAR secure database). The latter can be used in situations where there is a requirement to access more granular information (e.g., an emergency situation). The access to data regardless of the level of granularity is internally logged by TIPPERS so that attestation can be performed at any point.

In order to simulate a building within the NIWC Pacific campus with multiple rooms, the MCoE was divided into five areas: meeting space, visitor area, offices, machines, and kitchen. Each of these areas was represented as a *zone*. This zoning enables the definition of granular policies for different spaces (e.g., notify the administrator when a visitor moves to an area that is not a visitor or meeting space).

The graphical user interface (GUI) of the *Security Surveillance* application includes a heatmap showing the occupancy data. The current implementation includes the bird's eye view of NIWC Pacific topside, with simulated data to show the estimated occupancy of each building.




Figure 3. Mobility Center of Excellence room design.

The devices used in the deployment of TIPPERS in the MCoE include Bluetooth beacons, Wi-Fi APs, smart phones, Common Access Card (CAC) readers, cameras, microphones, and Raspberry Pis. Figure 3 shows the original design of the MCoE room along with the deployed sensors.

3.1.2 Demonstration use cases. Three use cases of interest for the Navy were showcased through the Security Surveillance application running on the TIPPERS deployment at MCoE.

- **Use case 1: exploratory occupancy analysis based on noisy and encrypted data.** The purpose is to highlight that surveillance tasks can still be done while preserving privacy. In this use case, the user accesses the heatmap first using noisy data and then requests permission to access more fine-grained data for one specific location where the occupancy is abnormally high.
- **Use case 2: automatic space policy enforcement in encrypted domain.** The purpose is to highlight that violations of policies can be automatically detected using the encrypted data—if a violation occurs further information can be de-encrypted. In this use case, the user defines a policy that says that visitors are not allowed in any space except for the visitor area. Then, a visitor carrying a smartphone moves to the office area and this event is detected by TIPPERS, which prompts the Security Surveillance application to display an alert.
- **Use case 3: breaking glass policy.** The purpose is to highlight that the Security Officer can access all the data if required in an emergency situation. In

this case, the user gets granted access to both real occupancy levels as well as trajectories of individuals. The interaction of the user with the system gets appropriately logged so attestation can be performed to analyze whether the access to such information was justified.

3.2 Instrumenting the Navy ship

As part of the deployment of TIPPERS in the assigned ship, the first step was to instrument the space with different IoT sensors. This required both the physical installation of the sensors and the deployment of a network infrastructure. With respect to sensors, the following list itemizes all of the equipment deployed on the Navy ship during TW 19:

- Wi-Fi APs (2);
- Bluetooth beacons (32);
- power outlet meters (6);
- Raspberry Pi (4);
- smart card reader (1);
- smartphones (30).

Each sailor participating in the testing of TIPPERS during TW 19 was issued a smartphone. WiFi APs, Bluetooth beacons, and smart card readers were used to passively locate people in the ship (through their assigned smartphones). In addition, smartphones were used to capture information about their integrated sensors (e.g., accelerometers and gyroscopes). Power outlet meters were used to capture information about energy utilization in the ship.

All the sensors listed above exchange data with the TIPPERS system via WiFi. The ship is equipped with a WiFi network infrastructure, based on the Navy's

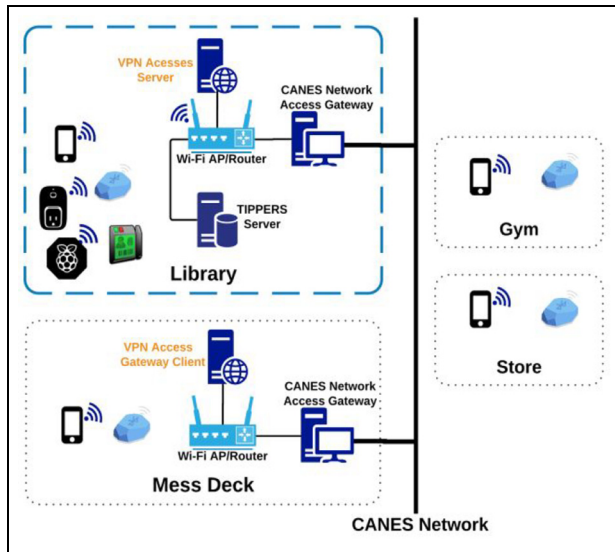


Figure 4. Sensor and network deployment. VPN: virtual private network; CANES: Consolidated Afloat Networks and Enterprise Services.

Consolidated Afloat Networks and Enterprise Services (CANES) infrastructure, which was leveraged for this task. Our design for the interconnection of sensors and the TIPPERS system by leveraging CANES is summarized in the schematic in Figure 4. The deployment leverages the CANES infrastructure by incorporating a site-to-site virtual private network (VPN) network on top of the CANES network. The purpose of this design is to minimize sensor configuration efforts. TIPPERS needs to collect data from various sensors for service provisioning. However, directly connecting a large number of heterogeneous sensors to the CANES network is not straightforward, since the CANES network has strict access control and firewall policies, which induce significant configuration efforts. On the other hand, building a completely independent network is also not practical. While easing the sensor configuration effort, it also poses the challenge of connecting the devices in different areas because the areas are far away from each other. A hybrid method was proposed, which connects sensors to the customized Wi-Fi network in each area and leverage CANES to route the data among these networks. In this configuration, each assigned space was instrumented with a CANES Network Access Gateway that was connected to the CANES network (as the assigned spaces were located far away). Then, in each space a Wi-Fi AP router that enabled the different sensors to communicate with the CANES Network Access Gateway was deployed.

A challenge that arises from the previous design is that data captured at each smartphone can only be transmitted to TIPPERS when the smartphone is connected to one of



Figure 5. TIPPERS mobile client.

the WiFi APs (located in the library and on the mess deck). Thus, data collected when the smartphone is located in other spaces (e.g., the gym or store) has to be stored in the device itself until the device moves to an area with WiFi connectivity. To this end, we developed a *TIPPERS mobile client application* for smartphones (see Figure 5) that handles this issue. The application continuously collected data from the smartphone's sensors (including information about the Bluetooth beacons around the device). At data collection time the application attempts to send such observations to TIPPERS through WiFi. If the smartphone is not connected to a WiFi AP, the underlying *data mule* technology in the TIPPERS client app stores those observations in an internal database on the smartphone. Then, when WiFi connectivity is re-established, the TIPPERS client application sends the stored data to TIPPERS using last in, first out (LIFO).

Sensor deployment in the ship was carried out without significant issues. The TIPPERS team was assigned four areas of the ship to instrument for the experiments: the mess deck, the library, the gym, and the store. The ship's library served as the focal point of the experiments and demonstrations and it was instrumented with one WiFi AP, six WeMo energy consumption monitors, four Raspberry Pis, a smart card reader, four Bluetooth beacons, and a MacBook Pro (where the TIPPERS system was installed and that also served as the demonstration station). The gym and store were both instrumented with two Bluetooth beacons. Finally, the mess deck was instrumented with 24

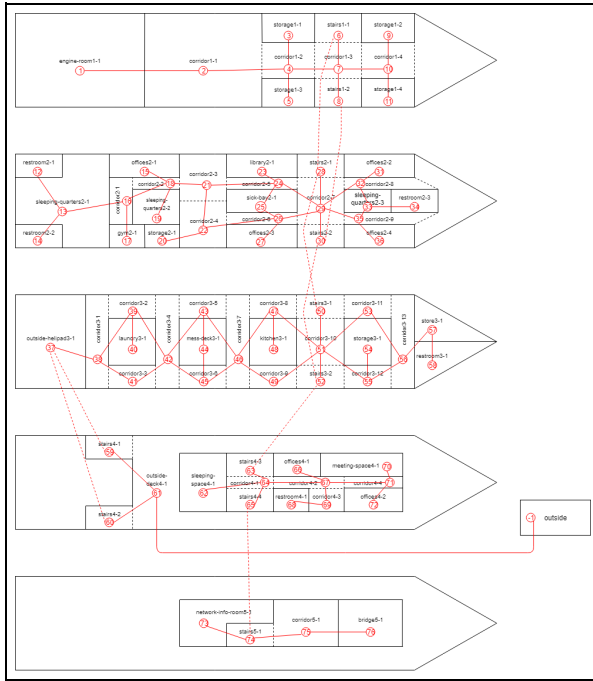


Figure 6. Definition of the ship's geographic information. (Color online only.)

beacons (one under each table), one WiFi AP, and one MacBook (working as a router).

Each of the Bluetooth beacons was placed and their positioning information was loaded into TIPPERS. The researchers used the iBeacon format with Major ID and Minor ID. Each Minor ID was a unique placement within the ship.

3.3 Configuration of TIPPERS

After the instrumentation of the ship, the next step was to configure the TIPPERS system. The TIPPERS system has been designed to be space-agnostic and to facilitate its deployment and configuration in different spaces with different underlying sensor infrastructures. TIPPERS is distributed as a docker container that includes all the required external libraries and software artifacts. To configure the system to the existing space, the first step is to insert, through the corresponding GUIs, geographical information of the ship (e.g., what types of rooms are there, their dimensions, etc.) along with information about people and their profiles. In total, as shown in Figure 6, more than 70 spaces in the ship were defined. Notice that given that the actual distribution of spaces in the destroyer is classified information, the figure shows a simulation based on declassified information for decommissioned US Navy destroyers. Also, the figure shows the adjacency of the

spaces with red arrows. This information was used to simulate synthetic trajectories of people on the ship, which is explained in the following section.

In addition, we configured and registered the devices to be used during the exercise. TIPPERS provides a GUI, which we call Portal, through which the administrator of the space (in this case the ship) can configure the system and register devices. The Portal is also the mechanism that users of the system use to register themselves along with their personal devices (e.g., their smartphones). Through the Portal, the users can access applications deployed in the space (e.g., to communicate with others or to see others' locations if allowed by individuals' privacy policies). During the exercise, sailors using the system register their smartphones (in this case provided by us) into the system. In a real deployment, users would access the system through the Portal and register themselves and their devices. This way, any device detected by the system that has not been registered (either by the administrator or the users) can be flagged as a potential threat.

Sensor wrappers were developed to enable TIPPERS to communicate with the specific sensors in the space. These wrappers encapsulate the low-level communication details (e.g., protocols, data formats) and send information to TIPPERS using its RESTful application programming interfaces (APIs) and JavaScript Object Notation (JSON) specification. Similarly, other software artifacts were created called *virtual sensors*, which translate low-level sensor data into higher-level semantically meaningful information that applications can leverage. For instance, one of the virtual sensors developed uses connectivity data from WiFi APs and beacons to create a notion of location of people in the space (i.e., it uses information about the coverage of the WiFi AP and the owner of the detected device to associate that person with the covered area at the time the connectivity event was captured). Other virtual sensors were also developed that use such location information to infer the occupancy of the different spaces.

3.4 Synthetic data generation/simulation

Given the limitation in the number of sensors that were deployed in the ship in the context of the Trident Warrior exercise, as well as the number of spaces available, a simulator tool was constructed to generate synthetic trajectories of people within the ship. In addition, the simulator is also used in the evaluation of the TIPPERS system, along with the different technologies included (such as secure storage or DP) prior to the deployment of TIPPERS in the ship. Along with that, given a description of a scenario, the simulator generates as output the trajectories of people as well as the connectivity events that WiFi APs would capture in such scenario. The description of a scenario that the simulator takes as input consists of the following.

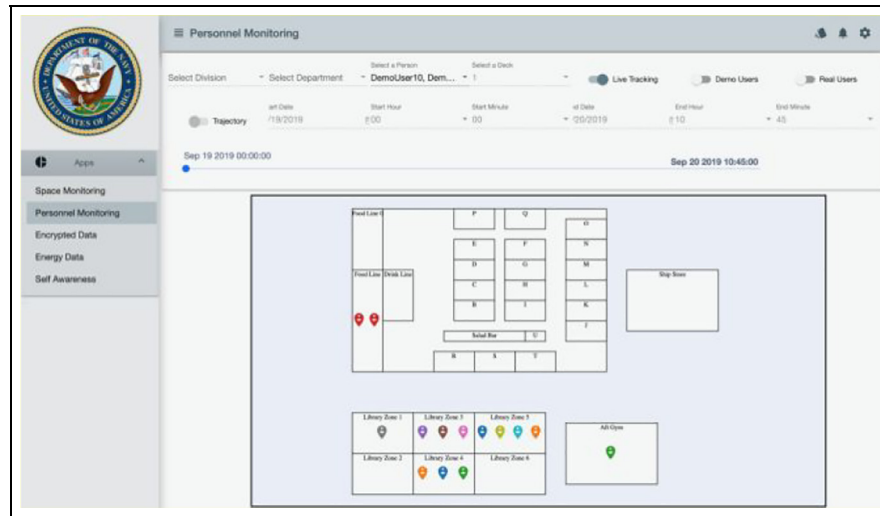


Figure 7. Command Board application.

- *Space definition:* the names and types of the different rooms in the space along with a directed graph. The vertices of the graph represent spaces in a scenario and the edges represent adjacency between spaces.
- *Event definition:* describes the events (e.g., cooking service hours) that occur in a space and are used as drivers of the simulation as people move in the space with the goal of attending events. Notice that some events are labeled periodic, meaning that they will occur on a periodic basis (e.g., every Monday).
- *People definition:* includes a template of the various individual profiles (e.g., sailors) in the scenario, including the events that each profile is likely to attend.

The information used to define the space, events, and people involved in the Trident Warrior exercise was provided by US Navy personnel. This included different profiles (e.g., Commanding Officer, Executive Officer, Command Master Chief, Officer of the Deck, Combat Information Officers, Main Propulsion personnel, etc.) as well as information about their schedules.

To generate the trajectories of people, the tool first uses the profiles input to generate the number of people per profile, which was included as a parameter and represents the expected numbers in a Navy destroyer. Then, for each individual, the tool assigns events that they can attend with some probability p (which is also part of the event definition and in this case was set up to 95% to represent in this context a sailor will most likely attend his/her daily duties).

For the events labeled periodic, a person attends an event when considered appropriate; this periodicity allows

the simulator to project the patterns that arise in the person's day. For additional noise in the simulation, people were given leisure breaks when they are not assigned to attend any event, and take restroom breaks throughout the day.

Synthetic trajectories were generated of the different personnel in the ship for the two weeks of the Trident Warrior exercise using the simulator tool. In total, more than 2.3 million connectivity data were created. These events were used to generate occupancy levels of the different spaces in the ship every 10 minutes. With this occupancy data, the Command Control application (explained in the next section) displayed occupancy at different temporal granularity (i.e., 10 minutes, 6 hours, 24 hours). In addition, the PGS virtual sensor was used to generate a stream of differentially private occupancy counts for those spaces. The goal was to evaluate how the noise included in the differentially private count would affect the utility of the data.

As differential private algorithms in general do not provide high utility when the counts are low (as adding noise, even if low, to small numbers decreases the utility further than to higher numbers), the focus was on the 6-hour and 24-hour cases. In addition, this data is also preloaded in the secure databases (see Appendix 1 (supplementary material)).

3.5 TIPPERS applications

Several applications were set up to showcase the benefits of the TIPPERS technology in the context of the Navy. All of the mobile applications were designed to run on both Android and iOS platforms. The Microsoft Visual Studio framework was utilized to develop the mobile

applications. The C# programming language and Xamarin were used as the language and application platform. Also, the main demonstration application (Command Board (CB)) was developed as a Web application that could be run on different computers used in the demonstration. All the developed software communicates with the TIPPERS system using RESTful API calls.

3.5.1 Command Board. The CB application provides situational awareness about the location of sailors and occupancy at different parts of the ship over time. Using the CB, the user (e.g., shipboard administrators) can monitor events and activities. In particular, Figure 7 shows a use case wherein an analyst uses the CB to visualize a heatmap of occupancy levels in the ship to explore anomalies, such as crowd formation or unexpected egress of people from certain regions. CB retrieves this data from TIPPERS that is internally generated by the PGS virtual sensor (see Figure 2) to ensure privacy of individuals. Detection of an anomaly over differentially private data can trigger further exploration in which the privacy might be traded off for improved accuracy of anomaly detection.

The CB also serves as a mechanism to inform the user about alerts such as unauthorized access of sensitive locations (by devices or users) and fall detection alerts (see Section 3.4.2). For the former, the CB can be used to denote certain regions as sensitive and the profiles of people who are authorized to visit them. This, along with the information that TIPPERS stores of devices registered to each user and people's location, can be used to trigger alerts when an unauthorized device/person enters a sensitive space. These alerts can result in decryption and display of trajectories of an individual, which by default is stored encrypted and inaccessible. Finally, the CB serves as an interface for messaging using the TIPPERS system as a backend and interacting with the TIPPERS messaging app (see Section 3.4.3).

3.5.2 Fall detection. As shown in Figure 8, the fall detection sensor application works in conjunction with other sensors and virtual sensors in TIPPERS to perform fall detection and verify fall conditions. If a fall is detected and verified, an alert is then generated by TIPPERS and sent to the CB for immediate action.

The fall detection application uses the popular root sum vector (SV) and threshold algorithm incorporated into many accelerometer-based fall detection systems.¹⁵ The software verifies falls detected by the client by cross-examination of other TIPPERS sensors, such as the WiFi AP, to confirm a fall has occurred and continues to provide more granular location information (i.e., sensor fusion). If the fall is detected near weather deck railings along with large, sustained accelerometer changes, a possible man overboard condition is computed.

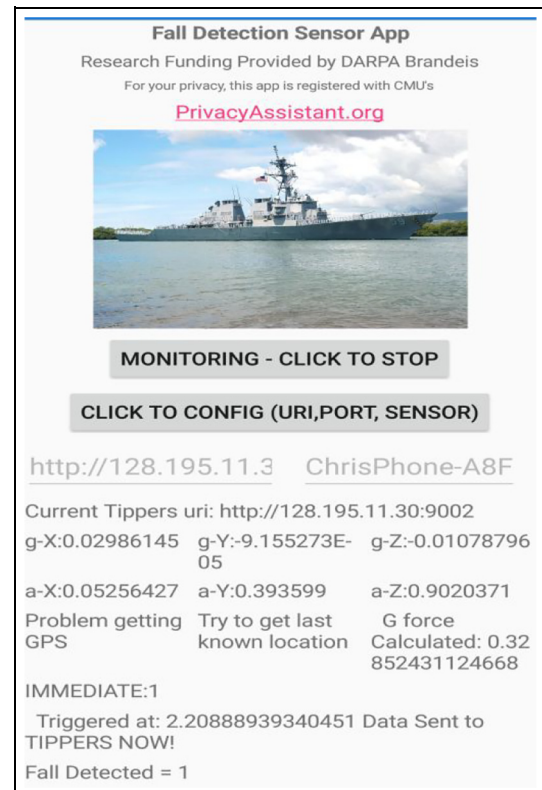


Figure 8. Fall detection application.

3.5.3 Messaging application. The TIPPERS messaging application supports point-to-point messaging, as well as point-to-multipoint (broadcast) messaging. Point-to-point messaging provides secure communications from the sender to the receiver. The point-to-multipoint broadcast functions as an intercom system. Such an ability could lead to replacement of the intercom in tactical environments where silence is required.

Users can compose messages both through the CB and the TIPPERS client application (see Figure 9). Each message is sent to TIPPERS, which stores it (encrypted) and then delivers it to the recipient's messaging application.

4. TIPPERS Trident Warrior 2019 use cases and experiments

In this section, we present the use case scenarios as they were designed and demonstrated during the TW 19 exercises and experiments.

4.1 Use case scenario descriptions

In order to test the various TIPPERS technologies and to provide real-world data for the simulations, 10 scenarios were tested. These scenarios came from several rounds of



Figure 9. Secure messaging application.

discussions with high-ranking Navy officers, a psychologist familiar with working with military members, and ex-military project members. The use cases were vetted and refined several times with the Navy before adoption by the project team. These scenarios (grounded in real-world Navy actuality) guided the technology testing as well as the technology testing during the TW 19 exercises.

The researchers divided these scenarios into two aggregates: mission critical and non-mission critical. Mission critical scenarios require immediate alerting and response. The non-mission critical scenarios provide planning and performance metrics that can be used for training and evaluation.

Scenario 1: privacy-preserving activity monitoring.

This scenario was used to understand sailor activity and movement within the ship and specifically focused on the mess deck. Reference points for time spent in line, time spent eating, and time moving in space were collected. Data for space utilization by the individual was collected as well. In addition, data on movement through the ship (e.g., store, mess deck, gym, and library) was collected for simulation purposes.

Scenario 2: command support. This scenario involved tasking sailors to carry out orders (e.g., movement from location to location) and provide command-level feedback on progress. This scenario also provided hands-on exposure to the commanders and sailors of the TIPPERS technologies.

Scenario 3: fall detection. This scenario was used to test the fall detection capabilities of the TIPPERS mobile client application (Figure 8) as well as the fall verification and alerting capabilities of the larger TIPPERS. While this scenario started as a man overboard scenario, it soon evolved into a more generalized fall detection use case based on feedback (e.g., stairwell dangers, shaft allies) from Navy personnel.

Scenario 4: physical security. In this scenario, the TIPPERS team tests physical security applications to




Figure 10. Experiment times |_2019-08-19. (Color online only.)

see whether the CAC reader triggers the expected alerts. This scenario was tested in the ship's library.

Scenario 5: space aggregation. This tests the ability of TIPPERS to provide aggregation and counts of sailors on the mess decks. The results show how many people use the mess deck, when they use it, where they sit, and their trajectories during usage.

Scenario 6: device management. This tests the ability of TIPPERS to capture specific characteristics of the supplied phones, both registered and non-registered. All phones used in this scenario are TIPPERS-supplied smartphones. Selected phones are unregistered and categorized as "rogue" cell phones.

Scenario 7: energy management. This monitors the energy consumption of TIPPERS equipment via the power outlet meter devices. Such devices, deployed in the ship's library, feed the TIPPERS system with data regarding energy consumption on individual electrical circuits.

Scenario 8: messaging and meeting scheduling. This scenario tests the TIPPERS secure messaging capability to the ship's company. The point-to-point (sailor to sailor) and the point-to-multipoint (intercom) features were tested (see Figure 9).

Scenario 9: activity self-awareness. This scenario has sailors carrying the TIPPERS mobile client while onboard. Then, sailors can use the CB application to monitor how many times they visited a specific area and/or how much time they spent there (e.g., how many hours they spent in the gym in the last month).

Scenario 10: SysAdmin and privacy. In this scenario, the TIPPERS team walks through the privacy leakage evaluation with Radio/IT personnel. The team solicits feedback on privacy as it relates to systems administration of TIPPERS.

4.2 Onboard data collection and experiments

The ship's crew participated in the data collection and experiments. Crew members were briefed on the nature of the data collection and provided informed consent (no individuals or individual identifiers were to be collected). Those that chose to participate were provided with a smartphone and were assigned a team member (observer) to visually verify the accuracy of the collected data. Each scenario (see Section 4.1) was tested within experiments and verified with an observer or demonstrated to Navy leadership (e.g., Scenarios 2 and 8). However, the mess deck space experiments (Scenarios 1 and 5) were by far the most time consuming and complex.

For the mess deck data collection, each sailor was approached for participation while they were in the mess line. If the sailor provided informed consent, they were instrumented with a phone and an observer logged their corroborative data. As the sailor moved through the mess line and the mess deck the data was observed and noted in a paper log while the TIPPERS mobile client application on their phones was sending sensor data back to TIPPERS.

An interesting point was voiced by many sailors with regard to the mess deck observations and data collection: privacy concerns. Almost every enlisted sailor (and many officers) voiced individual privacy concerns such as “are you tracking me?” or “how are you using this data?” to the TIPPERS team. The TIPPERS personnel explained that the data will be utilized to seed the system's simulators and provide space utilization metrics: no individual's identity will be attached to any movement data. This serves as a reminder that privacy within the military is a valued commodity and points out a fascinating dichotomy: data collection of individuals to study individual privacy.

5. Trident Warrior 2019 privacy study

An important part of the data collection during the exercise was the collection of data to infer the location of individuals. This data can be used to further compute various aggregated statistics (e.g., occupancy of spaces along time, average time spent by an individual in a particular location, average number of people an individual interacted with, etc.). Since the underlying data is private (i.e., location of individuals along time), the idea of an aggregated statistics leak of information about individual records is a precautionary concern. Well-known mechanisms can be applied, such as DP, to protect individual records, but knowledge on how to choose the appropriate privacy parameters is required. In this section, the different values of ϵ for a DP mechanism are estimated.

Assume a data table is displayed with user-location-time-table consisting of categories such as user, day, location, daytime, and time spent. The categories are used to

describe the amount of time the user has spent on each area per day. Consider, for instance, a commanding officer that observes aggregated averages on the time spent for each recorded location + daytime combination, stated as the following query:

```
SELECT day, location, daytime, AVG
(timespent)
FROM user-location-time-table
GROUP BY day, location, daytime;
```

The goal is to estimate how much a commander (treated here as an adversary) can learn about the particular time spent by a specific user. A quite strong attacker that may already have knowledge regarding the exact amount of time spent by other people who have been together with the victim at the same time in the same location can be assumed. This idea is motivated by the definition of DP, which is aimed to protect against such attackers.

In this experiment, the attacker first fixes a single victim out of n users. He/she computes the prior assumption of the victim's data based on the data of the other $n-1$ users (or only a certain fraction of these users). He/she tries to guess the victim's *spent times* (i.e., the amount of time spent in a particular area of the ship) based on the prior knowledge he/she has already learned, and on the aggregated statistics that depend on the victim's data. The researchers assume that the attacker wins even if he/she does not guess the spent time precisely, but with some precision. For example, if the attacker says that a user has been in a room for 17 minutes, but it actually was 17.5 minutes, the guess is still considered sufficiently correct.

There are n users that participate in the experiment. For each location + daytime + day combination, each user u_i has spent times distributed according to normal distribution $N(\mu_i, \sigma_i)$. The attacker predicts μ_i and σ_i based on the data of the other $n-1$ users.

Fixing some posterior probability t (e.g., $t = 0.9$), the researchers want to compute the precision r within which the attacker's guess stays with probability t . For example, if the actual time is x_0 , then with probability t the attacker's guess will be $x \in [x_0 - r, x_0 + r]$.

It can be assumed that an ϵ -DP mechanism is applied to the released average. In particular, the sensitivity of the AVG query with respect to attribute timespent is $1/n$, so, for example, the Laplace mechanism $\text{Lap}(\lambda)$ where $\lambda = 1/(n \cdot \epsilon)$ can be used.

Using existing results on relating DP to a guessing advantage (e.g., Pappachan et al.,⁸ and TIPPERS website⁹), if p_i is the prior guessing probability, then the posterior p'_i is bounded by the following:

$$p'_i \leq (1 + e^{R\epsilon} \cdot (1 - p_i)/p_i)^{-1},$$

where $R = \max_{x, x' \in X} d(x, x')$ is in this case the largest possible spent time. While normal distribution is unbounded,

it shows as $\Pr[x - \mu \leq a] = \text{erf}(a / (\sqrt{2} \cdot \sigma))$, where erf is the *error function*, so, for example, for $a = 3\sqrt{2} \cdot \sigma$ shown as $\Pr[x - \mu \leq a] = \text{erf}(3) \approx 0.9998$, which essentially covers the set of possible inputs. A smaller value of a can be taken to reduce the size of the exponent, but it also reduces the attacker's search space, so this parameter can be optimized to improve the upper bound on the guessing probability.

The team computes the following:

$$p_i = \Pr[x \leq x_0 - r] - \Pr[x \leq x_0 + r]$$

$$= \frac{1}{2} \left(\text{erf} \frac{|x_0 + r - \mu|}{\sqrt{2} \cdot \sigma} - \text{erf} \frac{|x_0 - r - \mu|}{\sqrt{2} \cdot \sigma} \right).$$

Then p'_i is computed from p_i and ε as described above. The experiments are performed for $r \in \{0.5, 1.0, 2.0, 3.0, 4.0, 5.0\}$ and $\varepsilon \in \{0.1, 0.5, 0.75, 1.0, 2.0, 3.0, 4.0\}$, computing the posterior p'_i . For each ε , the smallest r for which $p'_i \geq t$ is found.

Since the actual data of the users' movements on the ship collected during the exercise is classified and cannot be shared even for the privacy study, the researchers simulated the behavior of 150 users, such that the replicated data has the same statistical moments (means and standard deviations) as the actual data. This imitated data serves as the "real" data for the privacy study. Therefore, ε -DP is applied to this dataset and the adversary's success is computed.

The posterior guessing probabilities for five data samples had been estimated. The results are depicted in Figures 10–14. The number of spent times guessed is plotted with probability $\geq 90\%$ for different precisions, where the precisions are represented with different colors. The dark green color represents the roughest guess (± 5 minutes) and the dark red color the most precise guess (± 0.5 minutes). For $\varepsilon = 0.1$, only a few people are depicted in the bar, and this means that for the others the guessing precision was more than ± 5 minutes.

While the datasets are different, similar trends in these five plots are noticed. If $\varepsilon \geq 3$ is taken, then very few privacy guarantees are obtained, and each user's spent time may be guessed within 1 minute of precision. On the other hand, for $\varepsilon \leq 0.5$, the guessing precision ranges between 4 and 5 minutes, which is much better, considering that the actual spent times are on average 8–9 minutes in the given datasets. There are always several people for whom the guessing probability is large even for small ε , as their behavior is more predictable, but there are not too many such people. Since smaller ε means more noise in aggregated statistics, data utility also needs to be taken into account, which would be a separate study and depends on how the statistics are actually going to be applied.




Figure 11. Experiment times2_2019-08-19. (Color online only.)




Figure 12. Experiment times2_2019-08-21. (Color online only.)




Figure 13. Experiment times2_2019-08-22. (Color online only.)

Alternatively, weaker attackers could be taken, who do not know "everyone except the victim," but only some of the other users. In that case, it could be possible to get better privacy for larger values of ε . Modeling a particular attacker would require knowledge about the context, who the attacker is, and what he already knows. This remains outside the scope of this privacy study.




Figure 14. Experiment times3_2019-08-15. (Color online only.)

6. Conclusion

In this paper, the researchers presented the design and architecture of the DARPA-funded TIPPERS privacy-preserving pervasive computing platform and illustrated how the technological innovations were transitioned to enable new capabilities in situational awareness in a shipboard setting for the US Navy. TIPPERS was first deployed and validated in multiple real-world testbeds and then further deployed in the US Navy's TW 19 exercise on the Navy ship. Detailed scenarios were designed and implemented to cater to shipboard activities, test new technologies in a military environment, and facilitate a privacy study in tactical settings. The paper articulates how challenges in working with ship networks to execute IoT-based applications in a secure and privacy-aware manner were addressed. The PETs studied included policy aware data release, DP, and encrypted query processing. A privacy analysis on data collected during the exercise shows that DP technologies applied with the appropriate privacy parameter ϵ (i.e., ϵ -DP mechanisms) can be used to hide the precise time a sailor stays in a particular space while still offering some value for the analyst.

Lessons learned as a result of the technical deployment of the TIPPERS system in the TW 19 setting highlighted the important role of reliable real-time communications,^{16,17} privacy technologies, and tools for project management. Experiences gained from TW 19 are being used to design an enhanced deployment of TIPPERS that will be deployed in Trident Warrior 2020 while the ship is afloat.

A unique outcome of the privacy analysis from the TW 19 experience indicates that privacy within a military context is of paramount concern across the command spectrum. From senior officers down to the newly enlisted sailors, privacy concerns (especially in a ship setting) were many and varied. In many cases, privacy was critical to tactical operations as well as personnel security. In the future, minimally invasive architectures are critical to the

deployability of new and emerging IoT technologies that can create new levels of situational awareness while assuring the enforcement of privacy to individuals operating and working in the next generation of smart naval vessels.

Acknowledgements

The project team participated in TW 19. Trident Warrior is the US Navy's primary experimentation venue for Information Warfare (IW) technologies and adheres to the policies defined by Navy Warfare Development Command (NWDC) Fleet Experimentation (FLEX) program. The assistance from the Trident Warrior team and sailors aboard the ship was invaluable and greatly appreciated.


In addition, the team would like to acknowledge the following students for their support and assistance: Peeyush Gupta, Eun-Jeong Shin, Andrew Chio, Sameera Ghayyur, and Brianna Bowles.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Air Force, United States Navy, and DARPA.

Funding

This material is based upon work supported by the United States Air Force and DARPA (contract nos. FA8750-16-2-0021 and FA8750-16-C-0011) as well as the United States Navy and DARPA (contract nos. N66001-15-C-4065, N66001-15-C-4067, and N66001-15-C-4070). This work is partially supported by the NSF (grants 1527536 and 1545071).

ORCID iD

Christopher Davison  <https://orcid.org/0000-0003-0301-5745>

Supplemental material

Supplemental material for this article is available online.

References

1. Lee P, Shin E-J, Guralnik V, et al. Exploring privacy breaches and mitigation strategies of occupancy sensors in smart buildings. In: *TESCA* (eds Venkat P. Rangan, Nalini Venkatasubramanian, Maneesha Vinodini Ramesh and Serge Miranda), New York, NY, 13–14 November 2019, paper no. 7, pp.18–21. New York, NY: ACM Publisher.
2. Juels A. RFID security and privacy: a research survey. *IEEE J Select Area Commun* 2006; 24: 381–394.
3. Nayak R. *Radio Frequency Identification (RFID)*. Boca Raton: CRC Press, 2019. DOI: 10.1201/9781351238250.
4. DARPA Brandeis Program website. <https://www.darpa.mil/program/brandeis>. (accessed 21 August 2019)

5. Military News website. https://www.militarynews.com/news/navy-tests-experimental-technology-at-trident-warrior/article_55623050-dedf-11e9-8e21-4f20e90f315a.html (accessed 1 October 2019).
6. Mehrotra S, Kobsa A, Venkatasubramanian N, et al. TIPPERS: a privacy cognizant IoT environment. In: *2016 IEEE International conference on pervasive computing and communication workshops (PerCom workshops)* (eds Mohan Kumar and Aruna Seneviratne), Sydney, Australia, 14–18 March 2016, paper no. SP4, pp.1–6. Sydney, Australia: IEEE.
7. Yus R, Bouloukakakis G, Mehrotra S, et al. Abstracting interactions with IoT devices towards a semantic vision of smart spaces. In: *6th ACM international conference on systems for energy-efficient buildings, cities, and transportation (BuildSys 2019)* (ed. Hae Young Noh), Columbia University, NY, 13–14 November 2019, paper no. 10, pp.91–100. New York: ACM.
8. Pappachan P, Degeling M, Yus R, et al. Towards privacy-aware smart buildings: capturing, communicating, and enforcing privacy policies and preferences. In: *ICDCS workshops*, (eds Kisung Lee and Ling Liu), Atlanta, GA, 5–8 June 2017, paper number: Workshop Abstract 1, pp.193–198. Atlanta, GA: IEEE.
9. TIPPERS website. <https://tippersweb.ics.uci.edu/> (accessed 1 October 2019).
10. Boyle E, Gilboa N and Ishai Y. Function secret sharing. In: *EUROCRYPT* (ed. Tsonka Baicheva), Sofia, Bulgaria, 26–30 April 2015, paper no. 55, pp.337–367. Sofia: IACR.
11. Boyle E, Gilboa N and Ishai Y. Secure computation with preprocessing via function secret sharing. In: *TCC 2019* (ed. Dominique Schröder), University of Erlangen-Nuremberg Germany, 1–5 December 2019, paper no. 37, pp.341–371. Nuremberg, Germany: IACR.
12. Yao AC. Protocols for secure computations. In: *23rd annual symposium on foundations of computer science (sfcs 1982)* (ed. Michael J. Fischer), Chicago, IL, 3–5 November 1982, paper no. 21, pp.160–164. Chicago: IEEE.
13. Yao AC-C. How to generate and exchange secrets. In: *27th annual symposium on foundations of computer science (sfcs 1986)* (ed. Alok Aggarwal), Toronto, Canada, 27–29 October 1986, paper no. 19, pp.162–167. Toronto: IEEE.
14. Chen Y, Machanavajjhala A, Hay M, et al. PeGaSus: data-adaptive differentially private stream processing. In: *2017 ACM SIGSAC conference on computer and communications security (CCS '17)* (ed. Ahmad-Reza Sadeghi), Dallas, TX, 30 October–3 November 2017, paper no. 47, pp.1375–1388. Dallas: ACM.
15. Bagalà F, Becker C, Cappello A, et al. Evaluation of accelerometer-based fall detection algorithms on real-world falls. *PLoS ONE*. 2012; 7 (5): e37062.
16. Dustova G, Davison C, Hua D, et al. A model for first responder-academic collaboration. *Acad Exch Q* 2016; 20: 79–88.
17. Henderson LS, Stackman RW and Lindekilde R. The centrality of communication norm alignment, role clarity, and trust in global project teams. *Int J Proj Manag* 2016; 34: 1717–1730.
18. Gilboa N and Ishai Y. Distributed point functions and their applications. In: *advances in cryptology – EUROCRYPT 2014 lecture notes in computer science* (ed. Phong Q. Nguyen), Copenhagen, Denmark, 11–14 May 2014, pp.640–658. Copenhagen: DBLP.
19. Wang F, Yun C, Goldwasser S, et al. Splinter: practical private queries on public data. In: *NSDI 2017* (ed. Aditya Akella), Boston, MA, 27–29 March 2017, pp.299–313. Boston: USENIX.
20. Corrigan-Gibbs H, Boneh D and Mazières D. Riposte: an anonymous messaging system handling millions of users. In: *IEEE symposium on security and privacy* (ed. Sean Peisert), San Jose, CA, 18–20 May 2015, paper no. 16, pp.321–338. San Jose: IEEE.
21. Doerner J and Shelat A. Scaling ORAM for secure computation. In: *ACM conference on computer and communications security* (ed. Ahmad-Reza Sadeghi), Dallas, TX, 30 October–3 November 2017, paper no. 47, pp.523–535. Dallas: ACM.
22. Bunn P, Katz J, Kushilevitz E, et al. Efficient 3-party distributed ORAM. *IACR Cryptol. ePrint Arch* 2018; 2018: 706.
23. Ghayyur S, Chen Y, Yus R, et al. IoT-detective: analyzing IoT data under differential privacy. In: *SIGMOD conference* (ed. Bernstein), Houston, Texas, 10–15 June 2018, paper no. number 16, pp.1725–1728. Houston: ACM.

Author biographies

Dave Archer is a computer scientist and research leader at Galois, Inc.

Michael August is scientist in the Mobile Systems Engineering & Solutions Branch at NIWC Pacific.

Georgios Bouloukakakis is a postdoctoral researcher in the Donald Bren School of Information & Computer Sciences at the UCI.

Christopher Davison is an Associate Professor in the Center for Information and Communication Sciences at Ball State University and a researcher at the UCI.

Mamadou Diallo is a scientist in the Cybersecurity Engineering Division at NIWC Pacific.

Dhrubajyoti Ghosh is a PhD student in the Donald Bren School of Information & Computer Sciences at the UCI.

Christopher Graves is a scientist in the Information Assurance Division at NIWC Pacific.

Michael Hay is an Associate Professor in the Department of Computer Science at Colgate University.

Xi He is an Assistant Professor in the Cheriton School of Computer Science at the University of Waterloo.

Peeter Laud is a computer scientist and researcher at Cybernetica, AS.

Steve Lu is a mathematician and research leader at Stealth Software Technologies, Inc.

Ashwin Machanavajjhala is an Associate Professor in the Department of Computer Science at Dule University.

Sharad Mehrotra is Professor in the Donald Bren School of Information & Computer Sciences at the UCI.

Gerome Miklau is a Professor in the College of Information and Computer Sciences at the University of Massachusetts, Amherst.

Alisa Pankova is a computer scientist and researcher at Cybernetica, AS.

Shantanu Sharma is a postdoctoral researcher in the Donald Bren School of Information & Computer Sciences at the UCI.

Nalini Venkatasubramanian is a Professor in the Donald Bren School of Information & Computer Sciences at the UCI.

Guoxi Wang is a PhD student in the Donald Bren School of Information & Computer Sciences at the UCI.

Roberto Yus is a postdoctoral researcher in the Donald Bren School of Information & Computer Sciences at the UCI.