Defining the Internet of Devices: Privacy and Security Implications[1]

Georgia Institute of Technology Technical Report: GIT-GVU-14-01

Richard L. Rutledge, Aaron K. Massey, Annie I. Antón, and Peter Swire

Abstract—**What observers have called the Internet of Things (IoT) presents privacy and security challenges for contemporary society. The conceptual model of the IoT evolved rapidly from technologies used to track parts in industrial supply chain management to a diverse set of smart technologies. This rapid evolution has merged several conceptually distinct technologies into a single, difficult-to-define concept. A key difficulty is defining what constitutes a "thing." The term has been used to refer both to the things sensed, such as a star or the contents of a refrigerator, and to the things that do the sensing (devices). We argue that the Internet of Things is better conceptualized as an Internet of Devices (IoD) because devices, not things, act in a digital form and connect to the Internet. Along with the other requirements of an effective IoD, technologists and policy makers must develop standards, network protocols, identity management solutions, and governance for the IoD to address privacy and security challenges a priori rather than retrofitted after the fact. Privacy and security cannot easily be added to a system that is already deployed and established. In this paper, we define the IoT and the IoD and summarize the independent technologies from which they have evolved. We provide a five-stage general policy framework for evaluating privacy and security concerns in the IoD. Our framework seeks to provide a consistent approach to evaluating privacy and security concerns across all IoD technologies while remaining flexible enough to adapt to new technical developments.**

Index Terms—**Internet of Things, Internet of Devices, privacy, security**

I. Introduction

DEFINING the Internet of Things (IoT) can be challenging and confusing because colloquial definitions fail to accurately reflect the technologies in development and technical definitions are not easily mapped to real-world examples. What, exactly, is a thing and how does it relate to the Internet? As used colloquially and in the literature to date, things on the IoT may not be Internet-connected and may not even be electronic equipment. They might simply be every-day objects that are represented by networked data. Initially, things were tagged with machine-readable identification technologies, like advanced Electronic Product Codes (EPC) or Radio Frequency Identification (RFID) chips. However, IoT is now used to refer to sensors or smart objects that are themselves Internet-connected. Feki et al. estimate that 50 to 100 billion things will be connected to the Internet by 2020 [1].

This paper seeks to clarify the definition of the Internet of Things and provide a consistent set of terms for the various technical elements in the IoT. In particular, we introduce a consistent vocabulary that provides a way forward for technologists and policy makers seeking to mitigate threats to security and privacy that might result from IoT technologies.

An important part of this vocabulary is separating *things* from *devices*. Ryan Calo defines the IoT as referring "to the possibility of billions of devices—including everyday appliances such as your refrigerator—one day being networked and interactive" [2]. Although this definition accurately captures the IoT's excitement and promise, it does not identify the constituent technical elements in the IoT. In particular, a refrigerator is both a *thing* and a *device*. This dual role is not true of all objects that the IoT is considered to track. Consider that a refrigerator may track the groceries it stores so that it can automatically order replacements as needed, reducing the likelihood that individual consumers would run out of half and half for their morning coffee. In this case, the *things* being digitized, tracked, and made available for interaction are the contents of the refrigerator, but the *device* that makes this possible is the refrigerator itself. Some information about the refrigerator, such as the internal temperature, may also be digitized and made available, and in that case, the refrigerator would be both a thing and a device.

The Institute of Electrical and Electronics Engineers (IEEE) started a journal for research related to the Internet of Things in 2014, and their website defines the Internet of Things as follows:[2]

"The Internet of Things is a self-configuring and

R.L. Rutledge is a Ph.D. Student in the School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, USA, e-mail: rrutledge@gatech.edu

A.K. Massey is a Postdoctoral Fellow in the School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, USA, e-mail: akmassey@gatech.edu

A.I. Antón is a Professor in and Chair of the School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, USA, e-mail: aianton@cc.gatech.edu

P. Swire is the Nancy J. and Lawrence P. Huang Professor in the Law and Ethics program of the Scheller College of Business, Georgia Institute of Technology, Atlanta, GA, USA, e-mail: Peter.Swire@scheller.gatech.edu

[1] Presented at the 2014 Privacy Law Scholars Conference hosted by the George Washington University Law School in Washington, DC

[2] http://iot.ieee.org/about.html

> adaptive system consisting of networks of sensors and smart objects whose purpose is to interconnect "all" things, including every day and industrial objects, in such a way as to make them intelligent, programmable and more capable of interacting with humans."

This definition reflects an idealistic understanding of aspirations for the IoT rather than its current state. The IoT is not currently a fully self-configuring and adaptive system. The goal of connecting "all" things to the IoT, however, is further motivation for terminologically separating the things that are observed and the devices that observe them and exchange information with a network. There are an infinite number of things that will not themselves become part of the IoT. For example, stars are things in any ordinary use of the word, and telescopes can provide digital information about stars, but absent faster-than-light travel we will not make stars "intelligent, programmable, and more capable of interacting with humans."

The IoT began with an easy-to-define concept: a network for tracking things based entirely on easy identification. Radio-frequency identification (RFID) chips were added to otherwise mundane things so that RFID readers placed at important locations in a facility could identify them easily and efficiently. Since RFID readers can be omnidirectional, a network of fixed-place readers can provide complete facility coverage. RFID is used in many industries to track parts in warehouses, assembly lines, and retail stores. For example, if all the merchandise in a store had RFID tags, then checking out could be as simple as moving the shopping cart past an RFID reader all at once rather than scanning every item individually. As simple as RFID technologies are, they still change the security and privacy analysis from non-RFID enabled scenarios. The RFID tags that make checking out so easy could also make it easy for someone in the parking lot with an RFID reader to know exactly what you purchased as you walk to your car. Consider also the RFID passport issued by the United States government. RFID chips make accessing information on passports much more efficient for customs, but also expose users to potential security and privacy risks [3] such as skimming and eavesdropping by an adversary [4].

In this paper, we distinguish between things and devices as follows. Consider what happens if an RFID chip is removed from a piece of merchandise. The merchandise itself still exists as a thing, but it would no longer be connected to the IoT. This is a simple example of the need to disconnect things from devices. Alternatively, consider two systems for tracking cars as they travel through a city. In the first system, each license plate comes equipped with an RFID chip that can be read by an RFID reader at certain important intersections. The upper half of Figure 1 depicts this scenario with the RFID device reading the first license plate. In the second system, a high-speed camera capable of accurately interpreting license plates using image-processing algorithms reads each license plate. The lower half of Figure 1 depicts this scenario with the camera device reading the second license plate. Both systems are designed to track a thing category: license plates. However, the technologies used to do this are fundamentally different. In the first case, the license plate gained a new feature: the ability to broadcast its identity. In the second case, the license plate remains the same as it has for many years.

In this paper, we define *things* to be any object about which a device collects data or upon which a device operates. In turn, we define *devices* as the technologies that collect data from their environment, interact with their environment, and communicate through their network. Calo's definition refers to networking of "everyday appliances," and in this case, a refrigerator is both a thing and a device. Specifically, it is a thing with an embedded device.

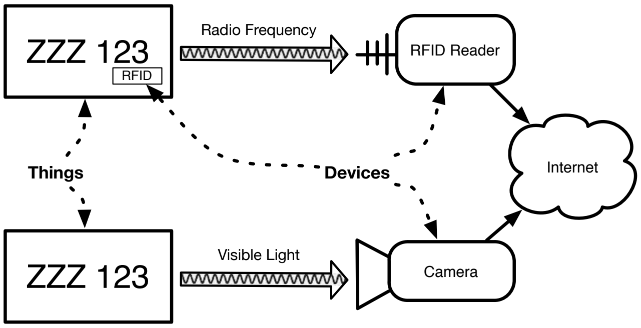

Figure 1. Differentiating Things and Devices in Systems for Tracking Cars by License Plates

We propose to substitute the term Internet of Devices (IoD) for the collection of technologies currently referred to as the IoT. Colloquial usage of the word 'thing' does not imply network communications. The term IoT no longer applies merely to devices that only provide identity information. Smarter devices capable of interacting with their environment are now commonly considered standard IoT devices. Consider thermostats that attempt to detect when homeowners are home and learn their travel patterns to improve the energy efficiency of heating and cooling the house. A traditional thermostat would not be a target for criminal activity, but a thermostat that learns when homeowners are absent might be. Google recently purchased Nest, a company that makes such a device, for $3.2 billion dollars [5]. An author for Wired described the purchase as something "that could finally bring the Internet of Things to life for all of us" [5]. Some Nest users also use Android phones, which are another Google product. When those users leave for home, their phones could tell their thermostats precisely when they left and, based on traffic, when to turn on the heat or air conditioning to ensure the home was prepared for their arrival. Although these technologies could be extremely convenient, there are security and privacy tradeoffs to adopting them.

Conceptualizations of the IoT are changing. The differentiation between things and devices matches ordinary English usage while providing conceptual clarity about what will be networked. That clarity will be important for analyzing the security and privacy issues that arise from increasingly smart devices that interact with their environment and possibly learn from the data they collect. Security and privacy professionals need a flexible framework for understanding and analyzing the different networks and devices that connect things on the IoD. In this paper, we categorize IoD networks, provide example devices and use cases for each type of network, and provide a framework for examining the resulting security and privacy implications. Our categorization covers five types of IoD devices:

1) **ID devices.** Simple identification-only devices that are physically attached to things (e.g. RFID)
2) **Remote sensors.** Devices that can learn to recognize and identify things remotely (e.g. cameras with product recognition software)
3) **Smart devices.** Sensors and articulators directly connected to (and potentially controlled through) the Internet. (e.g. home door locks that can be opened or locked using a mobile phone application)
4) **Application-specific computers.** General purpose computing devices connected to the Internet, but designed only for the purpose of running a particular application. (e.g. a mall kiosk)
5) **General-purpose computing devices.**

We begin our security and privacy analysis using the simplest possible framework, which only considers whether a device accepts inputs or generates data. We grow our framework progressively to allow for analysis of more complicated devices and situations resulting from the IoD. We also differentiate the IoD from other important security and privacy concepts. For example, IoD devices may generate huge amounts of data. Are the privacy and security implications of this data better thought of as Big Data concerns or IoD concerns? IoD devices may report data to or receive commands from third-party servers. Should the privacy and security concerns for these devices be thought of as Cloud Computing concerns or IoD concerns? IoD devices are deployed in a wide variety of contexts, including public infrastructure and wearable devices. Where do Ubiquitous Computing concerns end and IoD concerns begin?

Finally, we discuss how the pervasive availability of computing and networking exposes both government and private sector threats to individual privacy and security. The government routinely collects records of citizen activity, including employment records, income and property tax records, and voter rolls. What new records will the government collect when IoD devices are widely deployed? Although private industry cannot legally mandate data collection, companies regularly collect information as a by-product of a service paid for directly by customers (e.g., phone companies that collect sensitive records as a by-product of phone services.) or by providing services to customers and generating revenue through advertising (e.g., social networking sites that use various forms of advertising.). What new threats to privacy and security will be raised by business models for companies selling IoD devices and services?

The remainder of this paper is organized as follows. Section II introduces our terminology for the constituent elements of the IoD. In Section III, we review prior IoT research, focusing on security and privacy. In Section IV, we present our framework for examining the key security and privacy challenges of the IoD. In Section V, we outline a heuristic for conducting a privacy and security analysis for devices. In Section VI, we identify devices with key device attributes. Finally, we summarize our analysis in Section VII.

II. Terminology

The IoT conceptual model evolved considerably since its inception, but IoT terminology has not evolved with the model, fostering ambiguity and confusion. We now define the key terms that we employ for the remainder of this paper, beginning with *device* and *thing*. We provide parenthetical clarifications when discussing terms as used by other authors.

Device: A device is a combination of one or more components such as identifiers, sensors, or articulators (defined below) with a common control unit. If the device contains a sensor, then the composition must be uniquely identifiable. Similarly, if the device contains an articulator, then the composition must be addressable. An example device is an electric motor that reports its current speed and accepts commands for a new speed. Devices may, but do not have to, directly connect to the Internet. At least one component of a device must have some process for transmitting data to or receiving commands from the Internet. Consider a traffic sensor that collects data on the number of axles that pass over a particular section of highway. This device may not be directly connected to the Internet, but it is still considered a device if the data it collects is eventually made available either in a raw or aggregated form online.

Thing: A thing is any object about which a device collects data or upon which a device operates. For example, if a license plate scanner were installed and used to track the license plate numbers of every car passing through a particular intersection, then all of those cars would be things. This definition matches Privat's "extended things" [6], which we discuss in more detail in Section VI. We adapt Privat's notion of extended things because the key characteristic of "things" on the IoD is that they would otherwise be considered ordinary objects that do not by default produce data about themselves available on the Internet. When an ordinary object is targeted, tracked, or augmented to have a virtual existence,[3] it becomes a thing in the IoD. Moreover, if two devices were used to collect the same data set simply for redundancy purposes, this does not affect the number of things in the IoD.

Component: Components are the parts of a device that communicate over a network, collect data about the device's environment, affect state changes, or respond to identity requests. Sensors, articulators, and identifiers are all components.

Articulator: Articulators are components that accept commands through an IoD network and effect an appropriate change in physical or virtual device state. An articulator must be addressable on its network. Examples devices employing articulators include automated door locks and smart grid power switches. A less obvious example of a device with an

[3] There are some interesting parallels to Plato's Theory of Forms here, but they are perhaps outside the scope of this discussion.

articulator is a standalone GPS receiver. It receives commands from satellites and articulates by updating a local display.

Identifier: Identifiers are components that respond to identity requests. Identifiers may provide more than just identity information, but they can only provide information that they have been designed to provide. For example, an RFID is a device with an identifier component. It may be used to provide a unique identification number along with other information about the thing in which it has been embedded. If an RFID is embedded in a passport, it might include the name, address, and country of origin for the person to whom the passport belonged.

Sensor: Sensors are components that collect data about their environments and periodically transmit this data through an IoD network. Each sensor in each device must be uniquely distinguishable on their network. Example sensors include temperature and location sensors.

IoD Communications Protocol: A system of rules for data exchange across a network and between devices. Some devices may support multiple, simultaneous communications protocols and networks and route data between them. A smart phone may accept data over a Bluetooth protocol and forward it over a cellular protocol to a final Internet-based destination.

IoD Network: A set of devices that use a common IoD communications protocol to communicate with one another. IoD networks do not have to use communications protocols common to the Internet. Instead, they may choose to use a proprietary protocol for communications.

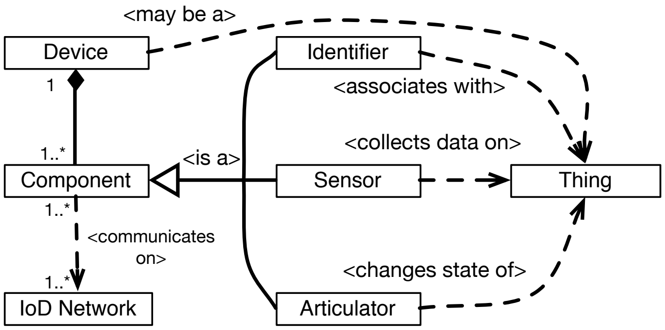

Figure 2. Relationship between Devices, Components, and Things

The relationship between these defined entities is depicted in Figure 2 using notation borrowed from the Unified Modeling Language class diagram. The figure shows the three device components and how each relates to a thing. Devices are composed of one or more components, each of which is an identifier, a sensor, or an articulator. Multiple components on a device may use an IoD communications protocol to communicate over an IoD network. A single component may also communicate over multiple IoD networks. For example, a typical smart phone can connect simultaneously to a cellular, Wi-Fi, and Bluetooth network and selectively route data between them.

Under these definitions, many devices are capable of rendering human beings as things. Although it may be strange to think of a human being as a thing, devices are often created with the explicit purpose of identifying, collecting data on, and interacting with humans. The Fitbit is a heath and fitness device that uses a variety of sensors to collect data on personal activity. Fitbits send data through a paired smart phone to company servers for analysis and later interaction. The data is collected only for the person carrying the Fitbit. Thus, the Fitbit is designed to treat the person as a thing. Carrying a Fitbit is not even the most invasive way that devices track human beings as things. In 2006, Applied Digital Solutions, Inc. sold over 1.7 million human-implantable RFID chips [7]. Human-implantable RFID chips wirelessly and automatically identify people for a variety of purposes, including medical records and payment systems. People like the convenience of being easily identifiable to computers.

People might not even be aware of the devices that identify them as things. Consider an RFID tag sown into the lining of a coat for store inventory management. The tag is an identifier device and the coat is the thing with which the tag is associated. However, if the purchaser is also known, then an association between the tag and the purchaser can be inferred by possession. That person is also now a thing. Due to the special properties of some remote sensors, a person does not have to be physically associated with a sensor to be a thing. A network of video cameras coupled with facial recognition software could track people's movements. In this example, a person associates with a device by simply and perhaps unwittingly walking into its range of view. We do not intend to de-humanize people by possibly categorizing them as things. Instead, we seek to accurately highlight the nature of the relationship between devices on the IoD and the people who carry them or are examined by them.

An object may be a device, a thing, both a device and a thing simultaneously, or neither a device nor a thing. Consider the smart home examples illustrated in Figure 3. In the diagram, boxes indicate objects. Rounded corners indicate that the object is also a component. Square corners with a single line border represent devices. Square corners with a double line border represent objects. Things are explicitly indicated with an arrow. A smart refrigerator containing an embedded RFID reader is on the left side of the figure. It can use the reader to enumerate its contents over an RF network. It contains exactly one RFID tagged carton. The tag is an identifier physically associated with the carton. This association makes the carton a thing. The refrigerator is connected to the Internet via Ethernet for product code lookup. The refrigerator is not a thing and remains only a device because no device, including itself, collects data about it, identifies it, or interacts with it. The embedded RFID reader collects data about the refrigerator's contents, not about the refrigerator itself. If the refrigerator's components (compressor, controller board, ice-maker) were RFID tagged and the reader capable of listing the refrigerator's constituent parts for a maintenance technician, then the refrigerator would be a thing in the IoD.

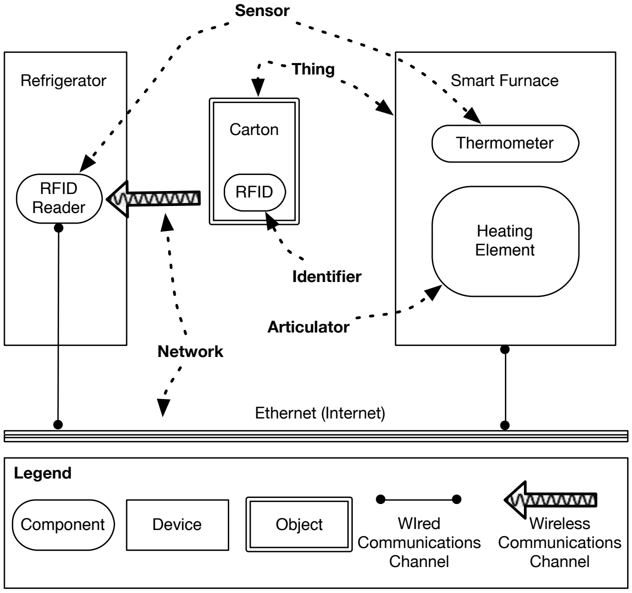

Figure 3. Smart Home Terminology Examples

A smart furnace appears on the right side of Figure 3. It contains two embedded components: a thermometer and a heating element. Although physically a part of the furnace, the thermometer detects the ambient temperature of the surrounding room and the furnace streams this data for external processing. The furnace also accepts commands to activate a heating element. The thermometer is insufficient to consider the furnace to be a thing because it does not collect data about the furnace. However, the external control of the heating element within the furnace does satisfy the definition and, thus, the furnace is both a device and a thing.

Our notions of "device" and "thing" differ from the terminology sometimes used in prior work. Some, perhaps most, technical research makes no such differentiation between devices and things, but we believe that highlighting this difference is helpful for both engineers and policy makers. Both disciplines are concerned with things and devices, but differ in approach. Engineers design and build devices that are required to observe or interact with things. The device is their objective. Policy makers seek to affect a characteristic of one or more things, such as the privacy of persons. One way of achieving that objective is to regulate the device. Thus engineers and policy makers may approach things and devices from opposing sides of a means-ends analysis.

III. Related Work

As computing devices became smaller, more power efficient, and cheaper to produce, they transformed early conceptualizations of an Internet of Tagged-Things into an Internet of Smart-Things, with a corresponding increase in exposure to security and privacy risks. In this section, we summarize and briefly examine the transition to the modern notion of the IoT. We first discuss radio frequency identification (RFID), where the term "Internet of Things" originated. Second, we discuss wireless sensor networks (WSNs), which mark the transition from devices that simply provide identity information to devices that can report on and interact with their environment. Third, we discuss IoT consumer applications in the smart home context. Fourth, we examine mobile computing and wearable computing. Finally, we examine an evolving Internet, which is changing from a network of computers and servers to include mobile and embedded devices. Throughout this section, when an author discusses an object, we parenthetically indicate the object's designation in our terminology defined in Section II. For example, when we write: "The author performed a threat analysis of RFID tags (identifier) in the Internet of Things (IoD)", the author originally wrote in terms of *tags* and the *IoT*. *Identifier* and *IoD* are the corresponding terms by our definitions.

A. Radio Frequency Identification

Kevin Ashton coined the phrase "Internet of Things" in a 1998 presentation to Procter and Gamble [8]. He said, "Adding radio-frequency identification and sensors to everyday objects will create an Internet of Things, and lay the foundations of a new age of machine perception [8]." In its simplest form, an RFID system (network) consists of an RFID tag (identifier), an RFID reader (sensor), and a computational device to process input from the reader. RFID tags can be attached to or embedded within the object to be tracked. Readers can discover tags in their immediate vicinity and query their identifiers. RFID tags generally contain only a small amount of data, and they sometimes only provide a unique identifier [9]. Servers maintain information about tagged things and index them on the unique RFID identifier, which allows readers to query servers for more information [9].

Ashton and others founded the MIT Auto-ID Center that envisioned objects (things) tagged, identified, and tracked by RFID [10]. The Auto-ID Center designed the Electronic Product Code (EPC) as a wireless, digital tag to replace the Universal Product Code (UPC), commonly called a bar code, in supply chain management [11]. Retailer RFID readers could calculate inventory levels and automatically order goods (things) as required. RFID tags are small, and may remain unnoticed by individuals who have purchased products that were tracked through the supply chain using RFID tags. Thus, individuals may be unaware of RFID tagged things in their possession, allowing them to be tracked without their knowledge or consent.

Early Internet regulation and policies significantly informed the European Union (EU) approach to the IoT with respect to privacy and security. In 2006, the EU Commission established IoT workshops, and in 2008 published a Staff Working Document [12]. After a comment period, the EU Commission published a plan[5] with 14 lines of action, three of which related to privacy and security:

Continuous monitoring of the privacy and the

[5] COM (2009) 278 final

protection of personal data questions. The Commission will develop guidelines on the operation of RFID applications in compliance with privacy and data protection policy and elaborate to include privacy and trust in a ubiquitous information society.

The "silence of the chips". The Commission will initiate a debate on the right of individuals to disconnect from their networked environment at will.

Identification of emerging risks. The Commission will provide a policy framework meeting the IoT challenges to trust, acceptance, and security.

In 2009, Rolf Weber summarized many EU concerns and actions taken to regulate the IoT [11]. At the time, IoT conceptual models were dominated by RFID and EPC. The not-for-profit association GS1 proposed the model developed at the MIT Auto-ID Center as the EPCglobal international standard.

IoT technologies and standards, such as EPCglobal, were initially based on Internet standards. Consider object naming and identification. EPCglobal has an Object Naming Service (ONS), which is based on the Internet's Domain Name Service (DNS) [11]. The Internet Corporation for Assigned Names and Numbers (ICANN) controls the Internet's DNS service. Weber criticizes ICANN's approach, claiming it lacks transparency and accountability [11]; he believes similar concerns will also apply to EPCglobal. For example, VeriSign currently operates the ONS directory service root node for EPCglobal and, as a result, has a great deal of practical influence over how EPCglobal operates. Weber concludes, "Since the IoT is not only a mere extension of today's Internet, [...] the development of decentralized architectures and the promotion of a shared network of multi-stakeholderism governance for the IoT is needed [11]."

At least two IoT (IoD) research groups proposed adapting the Platform for Privacy Preferences (P3P), an Internet standard for expressing privacy preferences [13], [14]. Tao and Peiran propose a P3P adaptation with three actors: an individual user, a service provider, and the 'third party' (a national or industrial supervisory party) [13]. They provide examples of information types and associated sensitivity levels, evaluation of user preferences by the service provider, and required authorities and responsibilities of the third party [13]. Ukil et al. also identified the individual data producer and data consumer as stakeholders in their negotiation-based privacy preservation technique [14]. They propose to extend the P3P XML-based schema to enable a Negotiation Module within the IoT to serve as an automated mediator between the individual and the service provider [14]. The Negotiation Module is governed by privacy policies that are, in turn, based upon privacy law [14]. Both Tao and Ukil use P3P to provide a basis for privacy in the IoT, but P3P has not been widely adopted because of concerns that limit its appeal [15]–[19]. In fact, some of the proposed P3P adaptations, like the automated negotiation modules that Ukil et al. propose, were not implemented in the original P3P specification due to implementation complexity, lack of interest from industry, and concerns that automated negotiation would not benefit consumers [19].

Similar to the above negotiation-based approach, Machara et al. propose to insert a Context Manager Middleware layer into the IoT (IoD) [20]. The context manager matches privacy requirements and guarantees encoded into data producer and consumer context contracts that are based upon meta-models. Rather than starting from a P3P baseline, the authors develop a model of an agreement between context producers and context consumers. Both the producer and consumer provide half-contracts that are matched at run-time by the context manger. The agreement is matched for one observable, i.e. the data to be read by the consumer [20]. If a match cannot be made between producer and consumer half-contracts, then the data is not accessible by the would-be consumer [20]. One of the advantages of this approach over the above P3P adaptations is in the ability to handle dynamic modifications to the producer and consumer contracts. Although these meta-models have been validated with a tool for checking model consistency, the Eclipse Modeling Framework, their complexity may be a significant obstacle to broad adoption, much like the above P3P negotiation.

The IoT inherits the existing Internet's security concerns multiplied by a greatly expanded scope and scale. An adversary would have little incentive to track an RFID encoded milk carton (thing). But an RFID encoded wallet (thing) linked to an individual, perhaps at point-of-sale, could be used to track its owner. Zhu et al. considered the security of connections between RFID tags (identifiers), readers (sensors), and backend systems such as the Object Name Service (ONS) [21]. They extend prior work on authenticated key exchange (AKE) in RFID systems (network) to handle mobile RFID readers (sensor), tag (identifier) corruption, reader (sensor) corruption, and multiple readers. The authors demonstrate the correctness of the proposed protocol and argue that it is more efficient than prior work in this area.

Instead of proposing new network architectural components to address security and privacy, Benjamin Khoo performed a threat analysis on a hypothetical GS1 EPCglobal RFID system exposed to the public domain [22]. Effectively, he modeled a future IoT as the existing EPC system without additional security protocols as safeguards. His analysis enumerated the following nine threats and effects [22]:

1) Rogue Reader: Read Confidential Data
2) Eavesdropping: Read Confidential Data
3) Relay Attack: Read and Write Confidential Data
4) Replay Attack: Read and Write Confidential Data
5) Tag Cloning: Read and Write Confidential Data
6) Tracking People: Read Confidential Data
7) Blocking: Denial-Of-Service
8) Jamming: Denial-Of-Service
9) Physical Tag Damage: Denial-Of-Service

Khoo emphasizes that the current technology represented by the EPC system was designed for supply chain management and is not sufficient for a public IoT (IoD). "RFID technology is great for tracking and keeping stock of items or animals but if this is applied to humans there have to be laws and regulation to govern its operation and strong enforcement or audit to ensure compliance as it can be so easily abused [22].” Here, humans are things, whereas an RFID is a device. He stresses that these issues must be pro-actively resolved before RFID technology can enable the pervasive and ubiquitous computing expectations of the IoT (IoD).

B. Wireless Sensor Networks

As a network of sensor devices, a WSN is a ready candidate for incorporation into what has been called the IoT. Extension of WSNs to the IoT, however raises the terminological confusion that has led us to define the IoD. The terminological confusion created by the word “thing” is understandable based on the history. The word “thing” applied initially to RFID chips, where the inventory item and the unique identifier were physically combined. When analysis expands to WSNs, then devices and things are entirely different. The technical and policy issues raised by the connection of sensors means that the analytic focus should be on the sensors, which are devices, rather than on the particular things they are sensing.

Wireless sensor networks (WSNs) predate IoT conceptual models. The Sound Surveillance System is one early example of a large-scale sensor network used by the US Department of Defense to track foreign submarines (things). In 1980, the Defense Advanced Research Projects Agency (DARPA) initiated the Distributed Sensor Networks (DSN) program [23]. Additional DARPA programs such as Sensor Information Technology further developed robust, ad hoc networking and distributed information processing [23]. At the same time, advances in microelectromechanical systems decreased the size, power consumption, and cost of sensors (devices) while simultaneously increasing their range. The addition of wireless communications technology enabled the transition from DSNs to WSNs.

A typical WSN is composed of a large number of self-contained, communicating sensor packages (devices) [24]. In the literature, these packages are often referred to as either sensor nodes or motes. The sensors are often densely distributed relative to their range in an ad-hoc manor and collaborate to provide observation data [24]. Each individual sensor node (device) may have minimal computational resources, but the aggregate network may have considerable computational capability [25]. Applications for WSNs include military, security, and environmental monitoring. Yick et al. categorize WSN applications as either Tracking or Monitoring [26]. Tracking targets include humans, animals, vehicles, and other objects (things). Monitoring targets include environmental conditions, patient health, factory automation, and other conditions (things).

Due to the ad-hoc nature of sensor node location, the network must be self-organizing [25]. Sensor nodes must be capable of discovering neighbor nodes and dynamically selecting data routes. Early WSN network topologies were predominately point-to-point and star designs that delivered data directly to a data collector (sink). Current WSN topologies operate on a mesh in which sensor nodes communicate with each other and collaboratively deliver observation data to the data sink [25]. These communications strategies provide resiliency in the complete system given individual sensor node failures and communications interference from physical obstacles.

The European Union’s IoT Architecture (IoT-A) project produced a proposed reference architecture that provides interoperability between RFID systems and WSNs. Within this context, Gessner et al. consider the requirements for object resolution functions [27]. The authors’ work adds dynamic and secure capabilities to object (device) name resolution as the WSN migrates into the IoT-A infrastructure. Gessner et al. propose requirements for Authorization, Authentication, Identity Management, Key Exchange and Management, and Trust and Reputation Architecture [27]. The authorization module is comprised of either Role-Based Access Control (RBAC) or Attribute-Based Access Control (ABAC). Authentication, Identity Management, and Key Exchange is based upon existing PKI principles. The Trust and Reputation component gathers behavioral information about entities in the IoT and assigns a trustworthiness rating to each entity that other entities can access to determine their level of interaction [27]. The authors do not specify this module in any detail. They indicate that fuzzy logic, Bayesian networks, analytical expressions, or bio-inspired algorithms could quantitatively measure trust. They further indicate that trust could be modeled as a Boolean value, a discreet range of values, or a continuous interval. As a requirements analysis effort, the first four modules are reasonably concrete. However, the trust module is too ambiguous to reason about its implementation and interoperability. A further weakness is that RFID tags lack the computational resources to meaningfully participate in elements of the proposed scheme, such as Authentication.

Incorporating WSN concepts into the IoT lead to early terminological issues that conflated the thing with the device. Privat challenged this assumption in his proposal to extend the Internet of Things to include mundane, un-communicating objects that attain thing-like properties due to the special remote tracking capabilities of some sensors [6]. He termed such objects sense-able things. For an example, we will consider a standard IoT transaction and compare it to a hypothetical extended IoT transaction.

The smart refrigerator has almost become a caricature, but will serve to illustrate. One smart refrigerator is equipped with a RFID reader. The consumer opens the refrigerator, removes a carton of milk, and empties the carton into a glass. He then discards the empty carton into a waste bin. Overnight, the refrigerator uses its RFID reader to enumerate its contents and fails to detect any milk. It places an order for a fresh carton from the Acme Corporation for delivery the next morning.

A second smart refrigerator is equipped with multiple interior cameras. The consumer opens the refrigerator, removes a carton of milk, and empties the carton into a glass. He then discards the empty carton into a waste bin. Overnight, the refrigerator uses it cameras to examine its contents. Its image detection and recognition software is unable to match with a milk container. It places an order for a fresh carton from the Acme Corporation for delivery the next morning.

The carton in the first refrigerator is a thing by the prior definition. It has been equipped with electronic communications, an RFID tag. The carton in the second refrigerator is just a standard inanimate object. Yet, the sensor in the refrigerator allows the object to be treated in much the same way. One could argue that the imagined carton is not quite a thing because the refrigerator cannot distinguish between two different milk cartons, and thus is not completely identifiable. Nevertheless, the resulting behavior is identical.

Similarly, an extensive network of automobile license plate scanners imbue un-ICT equipped cars with extended thing properties. And coupling facial recognition software with public surveillance cameras renders people as pseudo-things. Neither of these examples suffers from the quasi-identification of the refrigerator example. In each case, the object is uniquely identified.

Another potential concern with WSNs is that a security-compromised node (device) can reveal information from the complete system. Frequently, data is routed from node to node before eventually reaching the data sink. This ad hoc routing adds resiliency and robustness to the network, but exposes other sensor's data at a compromised node. Vladimir Oleshchuk surveyed secure algorithms to perform distributed computation. He defines and provides examples of Secure Multiparty Computations (SMC) [28]. SMCs allow collaborative computation without any party divulging its own input. As an example, consider Yao's Millionaire problem: How can two millionaires determine which is the richest without revealing their own net worth? In IoD terms, the same problem can be stated as determining which of two sensors reads the highest value without publishing either value. The author notes that generalized SMC solutions are impractical, but domain specific solutions can be suitable for constrained computing environments such as at the envisioned IoT (IoD).

Oleshchuk describes two further SMC algorithms [28]. The secure sum protocol would allow a set of sensors to compute the sum of their values without disclosing any of their values. The last SMC given determines the intersection of two sets without disclosing non-common elements to the other party. For example, consider a campus door access lock (device) and a university student desiring access. The lock has an authorization list to which it is programmed to permit access to the room. The student also has an authorization list associated with their identification card (identifier). The secure set intersection allows the lock to determine if the student should have access without the student disclosing his list of authorizations or the lock disclosing its list of permitted identifiers. Reliable and secure WSNs provide a strong technological basis for future smart home and ubiquitous computing development.

C. *Smart Homes and Offices*

The IoD promises to radically transform our homes and offices. Tom Coates's house in San Francisco provides a striking, if somewhat silly, example. Coates connected numerous sensors in his house to Twitter,[6] and the house tweets an appropriate statement when a sensor receives certain inputs. For example, Tom installed a motion sensor in his sitting room, so when the house detects someone sitting down, it tweets, "Pretty sure there's someone in the Sitting Room. @tomcoates is that you?" Coates has also installed moisture sensors for his house plants, temperature sensors in various locations, and light switch sensors so the house can tweet about the conditions of his ficus (thing), whether his air conditioner (thing) is operating, and when someone turns on the bedroom light (thing) [29]. Coates also uses web services to allow his @houseofcoates Twitter account to tweet that he's not home when he checks in somewhere else on Foursquare,[7] a location-based social network [29]. All of the devices Coates has used to allow his house to tweet are commercially available and relatively inexpensive.

The concept of a smart home is not new. In 1998, Georgia Tech began a project called the Aware Home Research Initiative [30]. This project's goal is to enable research into how a controlled home environment can improve health, wellbeing, entertainment, and sustainability for residents. In 2003, Cook et al. envisioned an agent-based approach for devices in a smart home to collect information on their physical environment, communicate this information to other devices, and make decisions based on this information regarding how to interact with their environment [31]. Algorithms like this have been in development for years, but the availability of sensors, like the ones Tom Coates uses to wire his home to Twitter, is cost prohibitive for most consumers. Most early research in smart homes focused on three areas: (1) improving healthcare, particularly elder care; (2) improving energy use through coordinated control of power-hungry appliances; and (3) improving daily life through entertainment and artificially intelligent convenient functions for the residents [32].

In parallel to smart home research, power systems researchers have explored how to build a smarter electrical grid. This smart grid research suggests that it may be possible for the power company to project with a great deal of accuracy which appliances a resident might be using and when they are using those appliances based entirely on the amount and signature of the power requirements of their house.

Homes are relatively well protected in many privacy legal regimes around the world. In the U.S., the Fourth Amendment explicitly protects homes from unwarranted searches. It remains unclear how the Fourth Amendment might apply to smart home or smart grid data collected by or stored on third-party servers. Current third party doctrine suggests that it will not receive as much protection as data stored inside the home.

Kanuparthi et al. addressed some security and privacy threats from the smart home with Physical Unclonable Functions (PUFs) [33]. PUFs are the hardware equivalent of a cryptographic one-way function and can be used in a

[6] The Twitter account can be found online at https://twitter.com/houseofcoates

[7] https://foursquare.com/

challenge-response protocol [33]. When presented with a challenge, an instance of a PUF (device) responds with a repeatable response. However, the response is unpredictable across different instances of the PUF, even if manufactured with the same process. Kanuparthi et al. foresee a vast number of smart, networked devices such as "medical implants, alarm clocks, wearable systems, automobiles, washing machines, traffic lights, and the energy grid" [33]. In our nomenclature, all of these objects are devices, but some may not be things. For example, if a wearable system collects data about the wearer, then the wearable is just a device, not a thing. To accomplish security and privacy in this environment, Kanuparthi proposes to integrate PUFs into IoT sensors and use PUFs for device identity management [33]. Existing cryptography can then provide secure channels through the network. Unlike a standard PUF, a sensor PUF accepts two inputs, a challenge and a physical quantity. For a given (challenge, quantity) pair the sensor PUF always produces the same response and the response is also unpredictable across other physical instances of the sensor PUF [33]. The principle limitation to Kanuparthi's approach is the reliability of current PUF manufacturing techniques and scalability to billions of devices.

Even with a cryptographically secure home IoT network in place, a tremendous amount of personal data will flow through a smart home. How can a resident verify who has access to their data? Mayer et al. approached this problem using data visualization [34]. They used a standard network protocol analyzer to inform an augmented reality user interface enabling the visualization of data streams both within the smart home and externally to remote services [34]. A visualization aide of this type may have lead to an earlier detection of an LG smart television leaking privacy data [35]. It remains to be seen whether this approach would scale to dozens or hundreds of home devices that could be connected to external services for reasons such as: checking for firmware updates, logging permitted biometric data, and ordering depleted pantry items. A device that is not complying with its privacy settings will be difficult to detect amongst the larger flow of valid traffic.

D. Wearable and Ubiquitous Computing

Edith Ramirez, the Chairwoman of the U.S. Federal Trade Commission (FTC) said in her opening remarks at the FTC Conference entitled "Internet of Things–Privacy and Security in a Connected World" that wearable healthcare devices are poised to revolutionize healthcare [36]. Wearable and ubiquitous computing (devices) may be poised to revolutionize more than just healthcare. Later in that same FTC Conference, Vint Cerf, the Chief Internet Evangelist at Google, said that Google Glass, an optical head-mounted display with a camera and microphone (composite), may one day allow a blind German speaker (thing) to have a conversation with a deaf American Sign Language speaker (thing). Though it is clear that wearable and ubiquitous computing devices will have an important affect on society in the near future, we are no closer to understanding the impact they will have on individual security and privacy.

One area where wearable and ubiquitous computing has already begun to affect society is in Location-Based Services (LBS). These devices introduce privacy concerns for IoD users because they could be misused for systematic mass surveillance. Recent development in mobile devices in terms of computational capacity, wireless connectivity, and geolocational devices enables portable access to location information. These devices include GPS satellite tracking, cellular tower triangulation, and WiFi fingerprinting and scanning. Any personal or wearable device that communicates regularly on standardize networks can also inadvertently regularly provide location information on the owner or user of the device.

Elkhodr et al. surveyed privacy risks in Android, Apple iOS, and Windows Mobile phones to illuminate the nature and scale of the problem [37]. Enabling LBS on these devices can deliver some compelling services to the end-user. Your phone can provide turn-by-turn directions to a desired location or identify the closest coffee shop. However, as Elkhodr reports, keeping that information private is more difficult than most users presume [37]. They refer to a report from Lookout, an anti-virus and security firm, that around 300,000 mobile phone applications have access to the user's personal data [37]. They also present the results of a joint study by Intel Labs, Penn State, and Duke University to monitor the behavior of a random sample of 30 out of the 358 most popular free applications for Android smart phones [37]. Of these 30, 15 of the applications were sending geographic location information to remote advertising servers. Seven of these applications even provided the phone's unique hardware identifier [37]. This would allow for data from one application to be matched up with data from any other application that also provides access to the unique hardware identifier.

In order to maintain the convenience of LBS without the corresponding privacy concerns, Liu et al. propose establishing a trusted middle-ware layer between the user and the service provider [38]. The phone's LBS request services through the middle-ware that relays the request to the service provider through a pseudonymous account [38]. Hu et al. also propose a middle-ware layer of software to provide emergency access to LBS data [39], but they make no claims regarding data privacy. Although Liu's approach has a few weaknesses, such as replacing a third-party service provider with a third-party middle-ware provider and the lack of a guarantee that a pseudonymous account will not be re-identified [38], it does highlight a current feature of most wearable and ubiquitous computing devices: they generally communicate through a single device.

Devices that can communicate over both WiFi and cellular communications networks can act as hubs and allow other devices that do not have WiFi or cellular connections to sync data to the Internet. Consider a Fitbit, which is a personal fitness tracker that must sync data to the Internet by way of some other device, such as a mobile phone. This model of a primary device upon which one or more satellite devices rely for communications is called a personal area network (PAN), and the IEEE is working on official protocols for PAN communications [40]. PANs offer a natural architecture for technical measures to protect privacy and security.

E. An Evolving Internet

A final independent set of technologies that are evolving into what is commonly considered the "Internet of Things" is the Internet itself. Enhancements to existing Internet protocols and capabilities may be made to accommodate the IoD. Researchers are looking at improving current Internet protocols and standards for IoT (IoD) adoption. Wang and Wen specified enhancements to the Domain Name System Security Extensions (DNSSEC) protocol [41]. The currently prevalent DNS has numerous security issues such as cache poisoning. DNSSEC adds public key cryptography to authenticate DNS database updates and verify the authenticity of DNS query results. Essentially, the server side is secured with public key infrastructure (PKI) so that the client can trust the server, but symmetrical processes are not provided. The authors propose the application of PKI to the client as well [41]. They do not provide a nomenclature for their enhancements, but herein their enhanced DNSSEC will be called DNS+. To prevent an attacker from bypassing DNS+ and using network addresses learned in some other fashion, DNS+ will not resolve things to physical network addresses, but rather to random pseudo- addresses unique to each communication session for public side access [41]. This scheme also requires a network gateway+ to map the pseudo-address to a physical address and to reject public side attempts for a direct connection to the physical address. The authors provide a security analysis to validate the proposed scheme. Nevertheless, several issues would impede practical applications of DNS+. Every consumer in the IoD would require a digital certificate, the routing protocols that underlie the current Internet would have to be revised to accommodate the gateway+, and every router would have to be able to determine a physical route to billions of things from a now randomized network address.

A combination of context aware access control and data transformations protect privacy in Huang et al.'s Privacy Preserved Access Control [42]. As with the user preference models previously discussed in subsection III-A, this model also entails a data producer (sensor), a data consumer, and the IoT (IoD) as a platform for securely sharing data. Raw sensor data from the producer is first transformed as per the producer's privacy settings. For example, individual data elements could be masked, stripped, or substituted with ambiguous values [42]. For data access, the authors describe a context aware, k-anonymity [43] policy and filter. They illustrate this point with an example of a producer/consumer pair who are colleagues and the data item is the individual's current location. When the producer is on-duty, the consumer is permitted to access the producer's exact location. When the producer is off-duty, a gridded location is returned satisfying k-anonymity [42]. Analyzing Huang's model with our definitions demonstrates the importance of the distinction we make between devices and things. In this model, the sensor has a privacy setting, not the thing being sensed. Hence, the model presumes a physical association between the device and the thing.

Evans and Eyers assert that access controls, such as RBAC and ABAC mentioned in subsection III-B, will not scale into the IoT (IoD) since these techniques require the naming of principles to be granted or denied access [44]. They maintain that in the highly dynamic environment of the IoT, ensuring consistent implementation of discretionary access would be impractical. They propose to use techniques from Information Flow Control (IFC) to directly label data packets with tagged values. This arrangement does presume the existence of a Trusted Computing Base (TCB) to mediate access to the data. By digitally signing the packet, the TCB can detect if the tags have been altered or removed. Tags should be assigned as soon as possible after the generation of the data, preferably by the sensor itself [44]. To overcome objections that tagging is too computationally expensive, the authors demonstrate an implementation of packet tagging on two low-cost, common embedded micro controllers. The author's approach is interesting, but it is difficult to overlook the requirement for a TCB. Also, a comprehensive ontology for the tagging of privacy related data would be difficult to achieve in advance. And once a tagging scheme was encoded into the embedded computational resources of the IoD devices, tag label management would likely be as difficult as access control principle name management.

The current management structure for the Internet may pose challenges for adoption as the IoT network. Weber considers national regulation, international agreement, and self-regulation as the appropriate legal source for IoT law [9]. He rejects national regulation as not meeting the IoT globalization requirements [9]. He acknowledges that neither international agreement nor self-regulation alone would be practical to implement and acceptable to preserving privacy [9]. He recommends a form of "co-regulation" in which government sets a general framework elaborated by the private sector [9]. Weber also notes the special difficulties in achieving globalization given the differing notions of privacy in various regions of the world [9].

IV. A Framework for Evaluating Security and Privacy in the IoD

A. Security and Privacy Analysis Matrix

Perhaps the simplest model for examining security and privacy on the Internet of Devices is a simple two-by-two matrix as shown in Figure 4. Stated as simply as possible, devices that accept input may have security concerns, and devices that produce output may have privacy concerns. Four types of devices exist in this model. Type 1 devices have both security and privacy concerns because they both accept inputs and produce outputs. Type 2 devices accept no inputs, but they still produce potentially many outputs. Therefore, Type 2 devices may have privacy concerns. Type 3 devices accept inputs, but produce no outputs. Type 3 devices may have security concerns, but they cannot have privacy concerns.

Type 4 devices accept no inputs or outputs, and they have no security or privacy concerns.

This framework is easy to interpret, but is it useful for evaluating security and privacy in IoD devices? Are there devices that cleanly fit into each of these categories? General-purpose computers are clearly Type 1 devices because they accept numerous inputs and are capable of producing even more outputs. Type 4 "devices" accept no inputs and produce no outputs. Technologies that fall into this category are unlikely to even be considered by society as technologies, like chairs, shoes, or hammers. These tools are so simple that they are considered to be everyday objects.

Examples of Type 2 and Type 3 devices are more challenging to identify. If a Type 2 device accepts no inputs but produces outputs, then it may correspond to a sensor as defined in Section II. Similarly, a Type 3 device, which accepts inputs but produces no outputs, may correspond to an articulator. In both cases, the devices must be "pure" to cleanly fit into these types. If an articulator produced even a small output, the device would need to be considered a Type 1 device. Similarly, if a Type 3 device accepted even a tiny input, it would actually be considered a Type 1 device. Although pure devices are extremely rare, they do exist. Security critical environments, such as air traffic control, commonly use bespoken unidirectional communications protocols. The Federal Aviation Administration standard Digital Altimeter Setting Instrument sensor transmits signals but does not have any circuitry to receive them. A Denial of Service could be performed by inducing noise on the transmission wire, but such an attack cannot impair the security of the sensor itself. Similar examples can be found for Type 3 devices. Smart locks for homes essentially just receive an input that tells the device to lock or unlock, but it produces no outputs.

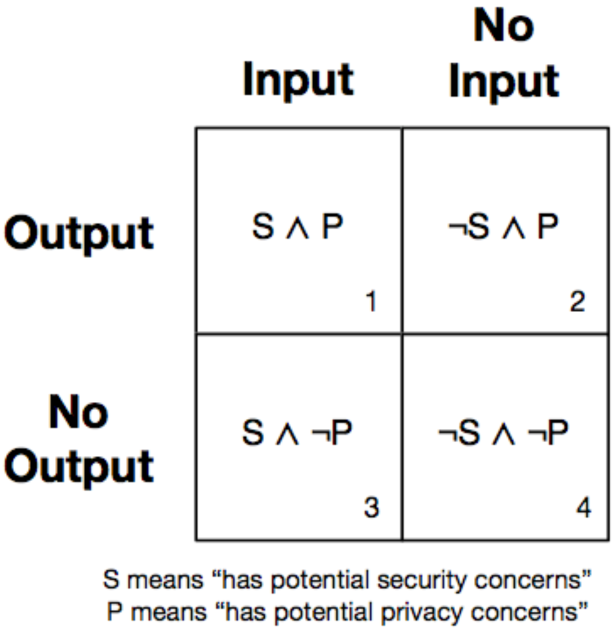

	Input	No Input
Output	S ∧ P 1	¬S ∧ P 2
No Output	S ∧ ¬P 3	¬S ∧ ¬P 4

S means "has potential security concerns"
P means "has potential privacy concerns"

Figure 4. Simple Analysis Matrix, showing four basic device types

Although Type 2 and 3 devices exist, they are the exception rather than the rule. Most devices run standard communications protocols that are inherently bidirectional. In order to both accept inputs and produce outputs, devices must include both sensors and articulators. If a device can be reduced to the sum of its component sensors and articulators, then we can perform an analysis as described in Figure 4.

To demonstrate that this relatively simple approach can still yield meaningful results, we apply the figure to two common, existing, and similar devices. A handheld, standalone GPS such as produced by Garmin or TomTom receives satellite data to calculate location, is addressable by virtual of being in range, and responds to data input by updating a user display. Since a handheld GPS cannot transmit back, it is a pure articulator. Applying Figure 4, we can conclude that although the GPS may possess a security risk, it is no threat to privacy. In comparison, we consider a smartphone with a built-in GPS. The smartphone also receives satellite data to calculate location, but does not contain a complete Geographic Information System (GIS) database. It would not be able to provide any further information if it did not also contain a sensor to relay the calculated location to the GPS mapping provider. In turn the mapping provider sends enough GIS data to the smartphone to update its display. The smartphone GPS is a composite of both sensor and articulator things. Applying Figure 4, we see that smartphones may contain threats to both security and privacy.

The framework as described thus far provides an oversimplification for actual devices, but it is useful as a starting point for our analysis. One way this is a simplification is the binary nature of the inputs accepted and outputs produced. A traffic counter designed to count the number of vehicles that pass over a section of a highway still has an output: the count of vehicles that have passed over the highway, possibly including time stamps for each vehicle. A smart traffic counter that could provide this data in real time would still have an extremely limited data collection process; it's still limited to a single section of highway and only capable of counting axles that pass over its sole input. However, these outputs might be available to other devices over a network and the ease of access fundamentally changes the nature of the traffic counter. Could a networked system of traffic counters identify reckless driving?

To address some of the limitations of our framework thus far, we may simply choose to examine devices based on their total number of inputs and outputs. This analysis allows us to create a continuous plot for devices rather than limiting our analysis to whether or not a device has any inputs. Devices with more inputs may be more of a security concern. Similarly, devices with more outputs could be more of a privacy concern. Figure 5 shows how this plot can be used to examine devices based on their total inputs and outputs. We might imagine that a city with networked parking meters would also want to install parking assistants, terminals posted on the street that could perform multiple parking functions. These terminals could accept additional inputs allowing them to serve users seeking to reserve a parking place at their destination. They could also allow a police officer to determine how long a particular vehicle has been parked outside the courthouse. Clearly, such a device poses both security and privacy concerns. Should the police be able to learn how long someone has been parked in a particular location? What if a combination of inputs exposed a bug that would allow anyone to learn that information, whether associated with law enforcement or not.

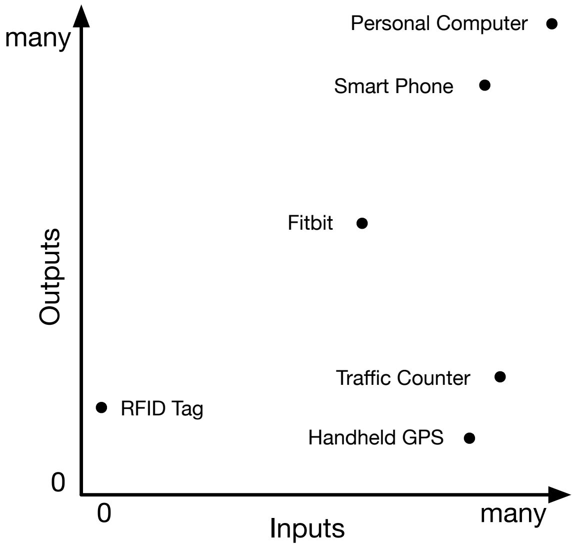

Figure 5. A continuous model for security and privacy concerns for devices on the Internet of Devices

The common use of standard bidirectional communications protocols means that many devices on the IoD will not be pure. This situation is exacerbated by the availability of low-cost, low-power general-purpose processors capable of running general-purpose operating systems. For example, Linux-based smart watches run an operating system that is strikingly similar to the operating system that powers a majority of web servers on the Internet. Even if a device is conceptually send-only, its underlying implementation may render it vulnerable to broad security threats. Security flaws in desktop operating systems are challenging to patch on a reasonable timetable. How much harder will it be to reliably update the security software installed on IoD devices that consumers do not even recognize as running an operating system?

Applications of our framework are intended to provide guidance rather than answers. Examining only inputs and outputs is a simplification, but it is still useful. Consider a hypothetical media device that accesses online media content for movies, music, and books. The device only sends a media title and receives the media content. On a data volume basis, the reception is several magnitudes larger that the transmission. Nevertheless, the device does warrant a significant privacy analysis due to the media titles and an implicit association of the device with a user account.

Inputs and outputs are not the only factors that may need to be examined. We may need to also distinguish between devices that communicate with people and devices that communicate with other devices. Devices that communicate with people could be considered the endpoints, the place where a security or privacy threat is actualized. Devices that communicate with other devices could be considered multipliers, which increase the impact a security or privacy threat might have once actualized. Our framework indirectly takes into account retention and archival of data transmitted by the device. A complete record of all media an individual has read or watched is a potential output, and thus greater risk to privacy than access to only the current media title.

Another factor to consider is whether a device is autonomous or dependent upon human interaction for its inputs or outputs. Devices that can communicate without human interaction may pose less of a threat to security and privacy since their dependence may allow for additional safeguards, such as authentication mechanisms, to be put in place prior to their communication. Devices that can communicate autonomously or automatically without human intervention may not allow for similar safeguards.

Extremely challenging analysis scenarios are easy to construct. Consider devices that accept many sensitive inputs, produce many sensitive outputs, and are able to communicate with either people or with other devices. They might communicate autonomously or with limited human input; for example, self-driving cars, autonomous drones, or citywide self-regulating traffic systems. Our framework begins to address the security and privacy concerns posed by these devices.

B. Device Categories

Our framework, which focuses on inputs and outputs, works well when used to analyze simple devices, but not all devices are simple. To examine how our framework applies to different devices, we introduce five device categories for the IoD. These device categories encompass a wide spectrum of types and computational capability as illustrated in Figure 6. We do not rigorously define each device category. Instead, we describe the general attributes of each category because the distinctions between categories are not easily made. Some devices could legitimately be analyzed from the perspective of more than one category. For each category, we provide example technologies that highlight core device characteristics representative of the category. In addition, we discuss briefly the applicability of our analysis framework, detailing the strengths and weaknesses of our approach.

1) **ID devices** are simple identification-only devices that are physically attached to things. These devices are only capable of responding to identify interrogation. Examples include RFID and Near Field Communication (NFC) tags. The development of this classification of devices inspired the early proposals for the IoT. The tag may be adhered to an object, or may even be integrally implanted into the construction of an object. Tags support detection by a separate interrogating sensor device and respond to queries with identity information. The response will include, at a minimum, a

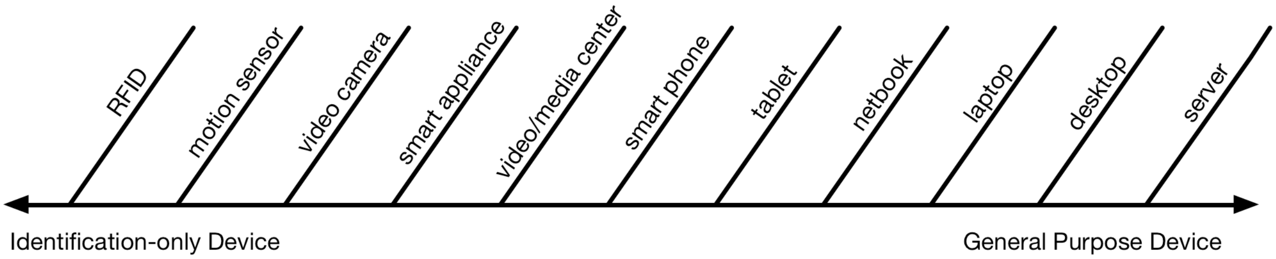

Figure 6. Spectrum of Devices in the IoD

classification of the attached object, such as a one-quart milk carton. Also, the limited range of RF and NFC allows the sensor to infer geo-location of the object. It must be near the sensor. These properties allow a RFID reader equipped smart refrigerator to determine if it contains a carton of milk and to not count a disposed carton in the trash bin. Additional information can be encoded in the identity response. Including the date of production would enable the smart refrigerator to recognize an expired carton and a serial number would enable it to identify a particular carton. The transmission of this data by the tag device and reception by the sensor renders the object a thing in our definition.

Simple identification-only devices are the closest devices to the ideal analysis outlined earlier in Figure 4. If the device has inputs, then it may have security concerns. If the device has outputs, then it may have privacy concerns. RFID devices accept as an input an RF signal that both indicates it should respond with its identifier and powers the device's ability to provide the response. This is a security concern because if the RFID device is damaged and unable to respond to this input, there will be no way for the reader to differentiate that device from a non-existent identifier. RFID devices also provide an output, the identification information. In the simplest case, this is information is built into the device when it is created, and the information cannot be updated after installed. Depending on the content of the information programed and the context in which it is accessed, this can be a privacy concern.

2) **Remote sensors** can learn to recognize and identify things remotely. These devices can render an object into a thing without physical attachment. A camera can remotely collect data about objects by recording electromagnetic waves. A sonar-capable device can perform similarly with audio waves. A smart refrigerator could be equipped with an array of internal cameras and product recognition software instead of a RFID reader. By periodically analyzing imagery from these cameras, this refrigerator could also determine whether it contains a carton of milk. Even though no device has been attached to the carton, the refrigerator is still able to collect data about the carton. Hence, the carton is still a thing and not just an object. This may seem like an overly sophisticated solution to design a refrigerator, but analogous situations already exist when security cameras are coupled with facial recognition software. This combination has been employed to detect suspicious individuals at sporting events [45]–[47] and renders these individuals as things.

Devices that use sensors to identify, recognize, and render objects as things also operate well with the basic analysis framework outlined in Figure 4. The sensors used to perform the recognition have an input, whether it is a photograph, a video, a scent, or some other potentially identifying data about a physical environment. This input is a potential security concern. If a license plate scanner is vandalized, perhaps by being covered in spray paint, then it cannot identify license plates. These devices also have outputs, which are privacy concerns. In contrast to the simple RFID device in the previous category, outputs from devices in this category may have a wide range of contextual privacy concerns. If a license plate has an RFID tag embedded in it, that tag may be read in contexts that are more revealing than the owner of the tag would prefer. If a license plate scanner uses a photography system to capture and read license plates, it may capture quite a bit more information than the just the license plate of the car. For example, it may capture an image of someone walking their dog on the sidewalk next to the car.

3) **Smart devices** are sensors and articulators directly connected to (and potentially controlled through) the Internet. These devices are constructed from dedicated hardware, operating system, and/or application software. They perform a narrow range of functions, and are not upgradable once installed. A smart-home owner could use a mobile phone application to open the garage door, unlock the entrance door, and turn on the household lights. The garage door opener, the entrance door lock, and the individual light fixtures are each examples of this category.

Devices with components that are directly connected to the Internet have security and privacy concerns that are not easily captured by a simple framework. A direct connection to the Internet is both a security and privacy concern simply because communication over Internet protocols requires both input and output. However, this description does not capture the myriad threats faced by devices directly connected to the Internet. If improperly mitigated, these threats might allow an attacker to remotely access the door lock to a house or office.

4) **Application-specific computers** are derived from general-purpose computing devices connected to the Internet, but designed only for the purpose of running a particular application. These devices may utilize general-purpose hardware, operating system, and/or application software. They perform comprehensive functions within an application domain, and are upgradable after installation. This large

category includes devices such as interactive, automated kiosks, smart phones, bank automated teller machines, and smart watches that run Linux-derived operating systems.

General-purpose computing devices that are designed to run a specific application face similar threats to security and privacy as smart devices. The key difference between them is that an application-specific computer may more easily be repurposed than a smart device. Consider a conference center kiosk that allows conference attendees to determine where sessions are located. This kiosk could be compromised by an attacker and turned into a node in a botnet. Worse, many kiosks deployed in this way have access to other computers on a trusted network. A compromised kiosk may allow an attacker access to other network resources and the information they contain. Organizations deploying and maintaining general-purpose computers intended to run a single application must maintain them as general-purpose computers rather than dumb terminals.

5) **General-purpose computing devices** must utilize general-purpose hardware, operating system, and/or application software. They perform a broad range of functions that are non-specific to any single application domain, and are upgradable at any time. Laptops, workstations, and servers can be firmly placed in this category. Other devices such as smart phones and tablet computers are challenging to classify since they possess attributes of both application-specific and general-purpose computers.

General-purpose computers have, as one might anticipate, quite a few security and privacy concerns. A simple examination of inputs and outputs is unlikely to suffice, and complete analysis is beyond the scope of this paper. However, the categorization of mobile devices, such as smart phones and tablets, as general purpose computing devices is an important consideration for the IoD. The non-computer look and feel of these devices may lead one to believe that they fall into an earlier category. Phones may even be thought of as everyday objects. It is critical that these devices are properly categorized and analyzed as general-purpose computers.

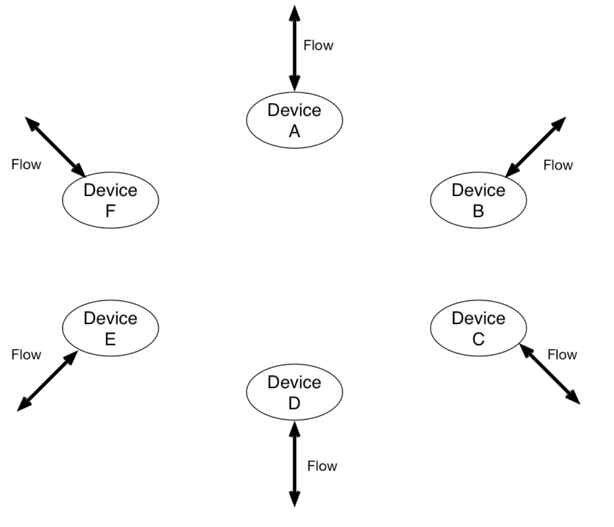

Figure 7. Stand-alone Devices

Neither the matrix nor the continuous model for security and privacy concerns directly account for the extent to which data is propagated on the network, although indirectly greater amounts of data propagated can be considered as outputs. Consider a Personal Area Network (PAN) as a composite device mediated by an Internet-enabled smartphone and containing biometric sensors for blood pressure, pulse rate, and body temperature. A straightforward application of the privacy/security matrix to the device indicates a high risk of privacy threat. However, if the sensor data were only used to provide the user with a status display on the smartphone and never sent the data upstream, then the scope of the sensor data is constrained to the phone. And the sensors do not constitute a privacy threat unless the security of the phone is compromised.

Black box analysis of device network communications may also complicate application of the models. Without access to the internal design details of a device, an analyst must resort to detection of transmitted and received data packets. However, detection may not be simple. For example, in November 2013 the BBC reported the discovery of a privacy breach committed by an LG smart television [48]. The complainant recognized that some form of tracking was taking place because the TV's user interface displayed targeted advertising. He possessed the tools and knowledge to investigate and found that the TV was sending channel selection information back to LG. Digging into the myriad of options, he found an opt-out configuration for "Collection of watching info," which he promptly turned off. Somewhat surprisingly, the TV continued to send channel selection information to LG in plain text, along with a flag to indicate the customer had opted out. Further, if a USB device is attached to the TV, it sends a list of all filenames found to LG.

This violation was found only due to the diligence of an IT professional. And LG could have evaded detection with only slightly more sophisticated technology or business models. If the data packet had been encrypted, it would have been more secure; even from the consumer. If the channel selection information been buffered and sent in bulk, it would have been more efficient; and less directly associated with channel selection. If LG has sold the collected information either to or in competition with Nielsen instead of selling targeted advertising, then this particular consumer would not have become suspicious.

The Federal Trade Commission (FTC) only engages after a consumer files a compliance complaint. How can technically proficient consumers detect non-compliance? In the LG television example above, the consumer used a simple form of flow analysis. Network packet monitoring software detected data packet flows that he did not expect to see. In part, this analysis was possible since the device was stand-alone and only required one flow analysis. As stand-alone devices accumulate, the number of flows to be analyzed increases linearly. In Figure 7, each of devices A – E can be analyzed separately. The addition of device F only requires analysis of a single additional flow.

Stand-alone devices are not, however, the objective of the

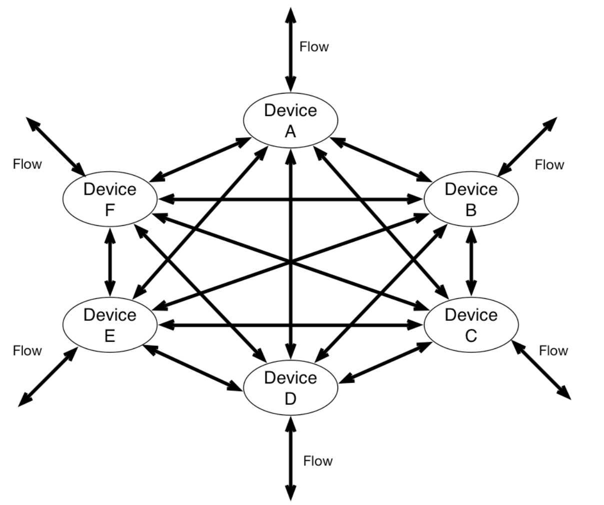

Figure 8. Collaborating Devices

IoD. In a recent press release, Samsung announced, "Samsung Smart Home's unique functionality enables users to control and manage their home devices through a single application by connecting personal and home devices—from refrigerators and washing machines to Smart TVs, digital cameras, smartphones and even the wearable device GALAXY Gear—through an integrated platform and server [49]." A flow diagram would look more like Figure 6 Figure 7 and the number of flows to be analyzed increases with the square of the number of devices.

In a fully connected mesh of collaborating devices where each device may be communicating with all other devices, the addition of a single device can have effects throughout the mesh. An additional device may require re-analysis of the entire en-meshed system rather than just the single device.

V. DIFFERENTIATING PRIVACY CONCERNS

The Internet of Devices clearly poses challenges for security and privacy, but it is not the only challenge for security and privacy. Mayer-Schönberger and Cukier discuss the massive amounts of data and metadata being created by devices in all levels of modern society as a Big Data privacy concern [50]. For devices that fall closer to the first two categories in the spectrum discussed in Section IV, the backend processing and analysis of data collected may best be thought of as a Cloud Computing Concern [51]. Autonomous devices, whether simple or complex, could justifiably be thought of as robots or drones, which also have a separate scholarship related to privacy [52]. Similarly, ubiquitous computing is another field in which scholars examine security and privacy concerns [53]. In this section, we discuss where each of these approaches to security and privacy may be applicable for devices and the IoD.

No bright lines separate Big Data security and privacy concerns from Internet of Things concerns. This lack of clarity is partially due to the challenge of defining both Big Data and the Internet of Things. Mayer-Schönberger and Cukier explicitly state that there is no rigorous definition for Big Data, and they instead choose to focus on the attributes of Big Data that are unique to Big Data [50]. For example, they say that "big data refers to things one can do at a large scale that cannot be done at a smaller one" [50]. Often, the ability to do things at a large scale will depend entirely upon data collected using IoD devices. Consider a large retailer, like Walmart, that uses RFID tracking on all of their merchandise in all of their stores. In this case, such a system allows the retailer to identify insights and opportunities that would not be possible without that scale.

Mayer-Schönberger and Cukier also describe big data as involving statistical calculations wherein the sample size is so large that it is effectively equal to the population size, allowing a transition from inferential to descriptive statistics [50]. IoD devices may enable these sorts of statistics. Consider the license plate tracking scenario discussed earlier. If license plate scanners are installed at every intersection in a city, it may be possible to track in real time all traffic at all times. City planners would no longer need predictive statistics to estimate traffic flows; they could simply use the actual number of vehicles. The IEEE's aspirational definition of the IoT claims that the purpose of the IoT is to "interconnect 'all' things," which is clearly related to the aspects of big data related to statistical calculations where n = all. If the "Database of Ruin" [54] is a consequence of Big Data, then a critical concern for the IoD is that it expands opportunities for growing the Database of Ruin.

If IoD simple devices, particularly devices closer to the first two categories in our spectrum, are intended to collect information on 'all' things, then they will need the support of Cloud Computing technologies. Discussions of cloud computing security and privacy concerns predate similar discussions regarding the IoD [51]. The continuous recording of data generated by IoD devices to a backend database substantially increases security and privacy risks. A house connected to a smart electrical grid can detect which devices are used by the residents, when those devices are used, and how long they are used. A power company collecting and processing this data on the cloud is in an excellent position to learn intimate details of individuals' lives. The amount of data that can be inferred by a smart meter is considerable, including identifying the program playing on the television [55]. Third party doctrine is a particular concern for privacy in cloud computing services [56], [57]. Cloud computing as also further exacerbated location-based jurisdiction issues for legal systems all over the world [58]. If IoD infrastructure is based on cloud computing technologies, then it will likely be beneficial to consider both approaches to examining security and privacy concerns.

Most cloud computing technologies are thought of as technologies, but the IoD emphasizes the extension of technology into spaces that are currently thought of as every day objects. As a result, cloud computing security and privacy concerns may not simply need to be considered in addition to IoD security and privacy concerns. In fact, the two areas may amplify one another. IoD devices we have discussed in this paper, like the smart refrigerator or the smart parking meter, may have serious implications for security and privacy specifically because they are not thought of as technologies. Bruce Schneier highlights the role that subtle social and technological cues inform trust and the implications these cues have on security and privacy concerns for the resulting socio-technical systems [59]. Technologies that are not thought of as technologies may prove to be riskier simply because people do not realize there are security and privacy risks or because people are more willing to forgo security and privacy in favor of convenience.

Autonomous devices and robots are another area where added convenience and utility may require a trade-off in security and privacy. Although IoD devices are not required to be robots in and of themselves, the aspirational view that IoD devices will be self-configuring, adaptive, intelligent, programmable, and more capable of interacting with humans is not dissimilar from the colloquial definition of a robot. In addition, current robots fill roles traditionally performed by people using common, everyday objects, which further suggests a shared set of security and privacy concerns between the IoD and robotics. Examples of these devices include iRobot's autonomous vacuum cleaner and Amazon's proposed drone- delivery system. Ryan Calo claims that robots raise privacy concerns "practically by definition" because they are able to "sense, process, and record the world around them" [60]. Certainly, Nest's learning thermostat fits this definition.

The IoD and robotics communities may overlap most in technologies that use artificially intelligent swarm-based algorithms. These technologies consist of simple devices that use basic interactions with their local environment or with one another to perform tasks leading towards emergent behaviors. These simple devices are closest to current IoD devices, and research in WSNs and self-configuring networks may naturally evolve to use swarm algorithms. Commercial applications of swarm-based robotics are not common yet, but it remains extremely promising and has a long history as a field of research [61]. Consider Google's driverless car project. Commutes would become shorter if every car on the highway participated in a swarm algorithm designed to mimic animal herding or bird flocking, but what are the privacy implications for those choosing not to participate?

Ubiquitous computing, often called ubicomp, is an umbrella concept that includes the colloquial understanding of the Internet of Things [53]. If the Internet of Things connects "all" devices, then ubicomp encompasses this concept and adds to it other concepts, like pervasive computing, haptic computing, distributed computing, and wearable computing. Researchers and technologists understand that ubicomp poses additional challenges to security and privacy [53], [62], [63]. The solutions and mitigations for those challenges may apply to IoD challenges as well.

VI. Summary

The conceptual model of the Internet of Things has evolved rapidly from a domain specific solution in supply chain management to a generalized platform for ubiquitous computing. Many open problems remain for technologists and policy analysts seeking to build, deploy, and regulate IoD devices, including privacy, security, standards, network protocols, identity management, and governance. Our paper provides three contributions that may address some of these open problems: (1) clarifying IoD definitions; (2) providing a framework for security and privacy analysis; and (3) providing guidance for where this analysis may need to be supplemented from other fields of research.

We began by addressing the confusing definitions for "things" in the Internet of Things. We introduce a concept for "devices," which refers to the technologies that collect data or interact with their environment, and differentiate them from "things," which refers to objects about which data is collected. Our clarification of "things" and "devices" includes a categorization of five types of IoD devices. These types are not rigidly defined, and they are best thought of as a spectrum of devices from simple identification-only devices to general purpose computing devices. Understanding the differences between these device types allows for a easier examination of security and privacy concerns. The more complex the device, the more complex the potential security and privacy concerns may be.

We also provide a simple framework for analyzing security and privacy concerns for devices on the IoD. Beginning with the simplest possible abstraction, we examine devices that accept inputs for security concerns and devices that produce outputs for privacy concerns. Although this simplification is not a perfectly representative abstraction, it can be useful in avoiding egregious errors of judgment. Furthermore, it is an extremely easy framework to apply to new devices.

Finally, we differentiated security and privacy concerns stemming from the IoD from security and privacy concerns that may best be examined under another context. In particular, we compared and contrasted concerns from the IoD, Big Data, Cloud Computing, Robotics, and Ubiquitous Computing. Each of these concepts has some overlap with technologies commonly considered to be part of the IoD, and understanding these areas of overlap is critical to properly resolving or mitigating security and privacy concerns for deployed systems.

The Internet of Devices will dramatically reshape the way we live and work. In some ways, the IoD is already here. The International Telecoms Union claims that at some point in 2014 cell phones will outnumber people [64]. The United Nations claims that more people have access to cell phones than toilets [65]. Consumers expect their every day objects to be smarter and more responsive than they did even a short time ago. A recent video of a 1-year-old attempting to treat a magazine like an iPad and finding it to be "broken" highlights how quickly this transition is taking place [66]. How soon will people who do not own a smart thermostat be as outmoded as people who do not have indoor plumbing? How quickly will cities with smart traffic analysis systems outnumber those that rely on upfront transportation planning? Most importantly, will technologists and policy analysts be prepared to examine the inevitable security and privacy concerns that arise when the IoD arrives?

References

[1] M. A. Feki, F. Kawsar, M. Boussard, and L. Trappeniers, “The Internet of Things: The Next Technological Revolution,” *Computer*, vol. 46, no. 2, pp. 24–25, 2013.

[2] M. R. Calo, “Digital Market Manipulation,” *University of Washington School of Law Research Paper*, no. 2013–27, 2013.

[3] R. Singel, “American Passports to Get Chipped,” *WIRED*, 2004. [Online]. Available: http://archive.wired.com/politics/security/news/2004/10/65412. [Accessed: 06-May-2014].

[4] M. Meingast, J. King, and D. K. Mulligan, “Embedded RFID and Everyday Things: A Case Study of the Security and Privacy Risks of the U.S. e-Passport,” presented at the RFID, 2007. IEEE International Conference on, 2007, pp. 7–14.

[5] M. Wohlsen, “What Google Really Gets Out of Buying Nest for $3.2 Billion,” *WIRED*, 14-Jan-2014. [Online]. Available: http://www.wired.com/2014/01/googles-3-billion-nest-buy-finally-make-internet-things-real-us/. [Accessed: 06-May-2014].

[6] G. Privat, “Extending the Internet of Things.,” *Communications & Strategies*, no. 87, 2012.

[7] I. Kerr, “The Internet of People? Reflections on the Future Regulation of Human-Implantable Radio Frequency Identification,” 2011.

[8] G. Santucci, “From Internet of Data to Internet of Things,” in *International Conference on Future Trends of the Internet*, 2009.

[9] R. H. Weber, “Internet of Things--New Security and Privacy Challenges,” *Computer Law & Security Review*, vol. 26, no. 1, pp. 23–30, 2010.

[10] H. Feng and W. Fu, “Study of Recent Development about Privacy and Security of the Internet of Things,” in *Web Information Systems and Mining (WISM), 2010 International Conference on*, 2010, vol. 2, pp. 91–95.

[11] R. H. Weber, “Internet of Things--Need for a New Legal Environment?,” *Computer law & security review*, vol. 25, no. 6, pp. 522–527, 2009.

[12] “Commission Staff Working Document, Future Networks and the Internet – Early Challenges regarding the ‘“Internet of Things,”’” *Samsung Electronics America*, 2008. [Online]. Available: http://ec.europa.eu/information_society/eeurope/i2010/docs/future_internet/swp_internet_things.pdf.

[13] H. Tao and W. Peiran, “Preference-Based Privacy Protection Mechanism for the Internet of Things,” in *Information Science and Engineering (ISISE), 2010 International Symposium on*, 2010, pp. 531–534.

[14] A. Ukil, S. Bandyopadhyay, J. Joseph, V. Banahatti, and S. Lodha, “Negotiation-based Privacy Preservation Scheme in Internet of Things Platform,” in *Proceedings of the First International Conference on Security of Internet of Things*, Kollam, India, 2012, pp. 75–84.

[15] T. Yu, N. Li, and A. I. Antón, “A Formal Semantics for P3P,” in *Proceedings of the 2004 workshop on Secure web service*, 2004, pp. 1–8.

[16] G. Hogben, “A Technical Analysis of Problems with P3P 1.0 and Possible Solutions,” in *Position paper, W3C Workshop on the Future of P3P*, 2002.

[17] G. Karjoth, M. Schunter, E. Van Herreweghen, and M. Waidner, “Amending P3P for Clearer Privacy Promises,” presented at the Database and Expert Systems Applications, 2003. Proceedings. 14th International Workshop on, 2003, pp. 445–449.

[18] I. K. Reay, P. Beatty, S. Dick, and J. Miller, “A Survey and Analysis of the P3P Protocol’s Agents, Adoption, Maintenance, and Future,” *Dependable and Secure Computing, IEEE Transactions on*, vol. 4, no. 2, pp. 151–164, 2007.

[19] L. F. Cranor, “Necessary but Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice,” *J. on Telecomm. & High Tech. L.*, vol. 10, p. 273, 2012.

[20] S. Machara, S. Chabridon, and C. Taconet, “Trust-Based Context Contract Models for the Internet of Things,” in *Ubiquitous Intelligence and Computing, 2013 IEEE 10th International Conference on and 10th International Conference on Autonomic and Trusted Computing (UIC/ATC)*, 2013, pp. 557–562.

[21] W. Zhu, J. Yu, and T. Wang, “A Security and Privacy Model for Mobile RFID Systems in the Internet of Things,” in *Communication Technology (ICCT), 2012 IEEE 14th International Conference on*, 2012, pp. 726–732.

[22] B. Khoo, “RFID as an Enabler of the Internet of Things: Issues of Security and Privacy,” in *Internet of Things (iThings/CPSCom), 2011 International Conference on and 4th International Conference on Cyber, Physical and Social Computing*, 2011, pp. 709–712.

[23] C.-Y. Chong and S. P. Kumar, “Sensor Networks: Evolution, Opportunities, and Challenges,” *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1247–1256, 2003.

[24] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, “Wireless Sensor Networks: A Survey,” *Computer networks*, vol. 38, no. 4, pp. 393–422, 2002.

[25] D. Culler, D. Estrin, and M. Srivastava, “Guest Editors’ Introduction: Overview of Sensor Networks,” *Computer*, vol. 37, no. 8, pp. 41–49, 2004.

[26] J. Yick, B. Mukherjee, and D. Ghosal, “Wireless Sensor Network Survey,” *Computer networks*, vol. 52, no. 12, pp. 2292–2330, 2008.

[27] D. Gessner, A. Olivereau, A. S. Segura, and A. Serbanati, “Trustworthy Infrastructure Services for a Secure and Privacy-Respecting Internet of Things,” in *Trust, Security and Privacy in Computing and Communications (TrustCom), 2012 IEEE 11th International Conference on*, 2012, pp. 998–1003.

[28] V. Oleshchuk, “Internet of Things and Privacy Preserving Technologies,” in *Wireless Communication, Vehicular Technology, Information Theory and Aerospace Electronic Systems Technology, 2009. Wireless VITAE 2009. 1st International Conference on*, 2009, pp. 336–340.

[29] R. Metz, “Home Tweet Home: A House with Its Own Voice on Twitter,” *MIT Technology Review*, 21-May-2013. .

[30] “About AHRI.” [Online]. Available: http://awarehome.imtc.gatech.edu/drupal/?q=content/about-ahri. [Accessed: 12-May-2014].

[31] D. J. Cook, M. Youngblood, I. Heierman, E.O., K. Gopalratnam, S. Rao, A. Litvin, and F. Khawaja, “Mavhome: An Agent-Based Smart Home,” in *Pervasive Computing and Communications, 2003. (PerCom 2003). Proceedings of the First IEEE International Conference on*, 2003, pp. 521–524.

[32] M. Chan, D. Estéve, C. Escriba, and E. Campo, “A Review of Smart Homes-Present State and Future Challenges,” *Computer Methods and Programs in Biomedicine*, vol. 91, no. 1, pp. 55–81, Jul. 2008.

[33] A. Kanuparthi, R. Karri, and S. Addepalli, “Hardware and Embedded Security in the Context of Internet of Things,” in *Proceedings of the 2013 ACM Workshop on Security, Privacy &#38; Dependability for Cyber Vehicles*, Berlin, Germany, 2013, pp. 61–64.

[34] S. Mayer, C. Beckel, B. Scheidegger, C. Barthels, and G. Šoř̈os, “Demo: Uncovering Device Whispers in Smart Homes,” in *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia*, Ulm, Germany, 2012, pp. 56:1–56:3.

[35] L. Kelion, “Lg Investigates ‘Spying’ Smart Tvs,” *BBC News*, 2013. [Online]. Available: http://www.bbc.com/news/technology-25018225. [Accessed: 17-Apr-2014].

[36] Federal Trade Commission, “Internet of Things - Privacy and Security in a Connected World.” [Online]. Available: http://www.ftc.gov/news-events/events-calendar/2013/11/internet-things-privacy-security-connected-world. [Accessed: 06-May-2014].

[37] M. Elkhodr, S. Shahrestani, and H. Cheung, “A Review of Mobile Location Privacy in the Internet of Things,” in *ICT and Knowledge Engineering (ICT Knowledge Engineering), 2012 10th International Conference on*, 2012, pp. 266–272.

[38] J. Liu, X. Hu, Z. Wei, D. Jia, and C. Song, “Location Privacy Protect Model Based on Positioning Middleware among the Internet of Things,” in *Computer Science and Electronics Engineering (ICCSEE), 2012 International Conference on*, 2012, vol. 1, pp. 288–291.

[39] C. Hu, J. Zhang, and Q. Wen, “An Identity-Based Personal Location System with Protected Privacy in IoT,” in *Broadband Network and Multimedia Technology (IC-BNMT), 2011 4th IEEE International Conference on*, 2011, pp. 192–195.

[40] A. Lo, W. Lu, M. Jacobsson, V. Prasad, and I. Niemegeers, “Personal Networks: An Overlay Network of Wireless Personal Area Networks and 3G Networks,” in *Mobile and Ubiquitous Systems - Workshops, 2006. 3rd Annual International Conference on*, 2006, pp. 1–8.

[41] Y. Wang and Q. Wen, “A Privacy Enhanced DNS Scheme for the Internet of Things,” in *Communication Technology and Application (ICCTA 2011), IET International Conference on*, 2011, pp. 699–702.

[42] X. Huang, R. Fu, B. Chen, T. Zhang, and A. W. Roscoe, “User Interactive Internet of Things Privacy Preserved Access Control,” in *Internet Technology And Secured Transactions, 2012 International Conference for*, 2012, pp. 597–602.

[43] L. Sweeney, “K-Anonymity: A Model for Protecting Privacy,” *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, Oct. 2002.

[44] D. Evans and D. M. Eyers, “Efficient Data Tagging for Managing Privacy in the Internet of Things,” in *Green Computing and Communications (GreenCom), 2012 IEEE International Conference on*, 2012, pp. 244–248.

[45] T. Perry, “Hockey Fans to Test Facial Recognition Technology - IEEE Spectrum.” [Online]. Available: http://spectrum.ieee.org/tech-talk/computing/software/hockey-fans-to-test-facial-recognition-technology. [Accessed: 12-May-2014].

[46] P. Rolfe, “Facial recognition technology to help combat sport troublemakers,” *HeraldSun*. [Online]. Available: http://www.heraldsun.com.au/news/law-order/facial-recognition-technology-to-help-combat-sport-troublemakers/story-fni0fee2-1226828714068. [Accessed: 12-May-2014].

[47] R. King, “U.S. testing crowd-scanning facial recognition system.” [Online]. Available: http://www.biometricupdate.com/201309/u-s-testing-crowd-scanning-facial-recognition-system. [Accessed: 12-May-2014].

[48] L. Kelion, “LG investigates Smart TV `unauthorized spying’ claim,” *BBC News*, 20-Nov-2013. [Online]. Available: http://www.bbc.com/news/technology-25018225.

[49] Samsung, “SAMSUNG Unveils New Era of Smart Home at CES 2014,” *Samsung Electronics America*. [Online]. Available: http://www.samsung.com/us/news/22331. [Accessed: 17-Apr-2014].

[50] V. Mayer-Schönberger and K. Cukier, *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. Eamon Dolan/Houghton Mifflin Harcourt, 2013.

[51] E. Qin, Y. L. C. Zhang, and L. Huang, “LNCS 8017 - Cloud Computing and the Internet of Things: Technology Innovation in Automobile Service,” pp. 1–8, Jun. 2013.

[52] R. Calo, “People Can Be So Fake: A New Dimension to Privacy and Technology Scholarship,” *Penn State Law Review*, vol. 114, no. 3, 2010.

[53] L. J. Camp and K. Connelly, “Digital Privacy: Theory, Technologies and Practices,” A. Acquisti, S. D. C. di Vimercati, S. Gritzalis, and C. Lambrinoudakis, Eds. Taylor & Frances, New York, NY, 2007.

[54] P. Ohm, “Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization.,” *UCLA Law Review*, vol. 57, no. 6, 2010.

[55] U. Greveler, B. Justus, and D. Loehr, “Multimedia Content Identification Through Smart Meter Power Usage Profiles,” *Computers, Privacy and Data Protection*, 2012.

[56] J. Harper, “Reforming Fourth Amendment Privacy Doctrine,” *American University Law Review*, vol. 57, p. 1381, 2008.

[57] C. Soghoian, “Caught in the Cloud: Privacy, Encryption, and Government Back Doors in the Web 2.0 Era,” *J. on Telecomm. and High Tech. L.*, vol. 359, 2009.

[58] D. R. Desai, “Beyond Location: Data Security in the 21st Century,” *Communications of the ACM*, vol. 56, Jan. 2013.

[59] B. Schneier, *Liars and Outliers: Enabling the Trust that Society Needs to Thrive*. John Wiley & Sons, 2012.

[60] R. Calo, “Robot Ethics: The Ethical and Social Implications of Robotics (Intelligent Robotics and Autonomous Agents series),” P. Lin, G. Bekey, and K. Abney, Eds. The MIT Press, 2011.

[61] T. Balch and R. Arkin, “Communication in reactive multiagent robotic systems,” *Autonomous Robots*, vol. 1, no. 1, pp. 27–52, 1994.

[62] B. A. Price, K. Adam, and B. Nuseibeh, “Keeping ubiquitous computing to yourself: A practical model for user control of privacy,” *International Journal of Human-Computer Studies*, vol. 63, no. 1–2, pp. 228–253, 2005.

[63] L. Strahilevitz, “Reputation Nation: Law in an Era of Ubiquitous Personal Information,” *Northwestern University Law Review*, vol. 102, Oct. 2008.

[64] “2014: Mobiles ‘to outnumber people,’” *BBC News*. [Online]. Available: http://www.bbc.co.uk/news/technology-22464368. [Accessed: 12-May-2014].

[65] Y. Wang, “More People Have Cell Phones Than Toilets, U.N. Study Shows,” *TIME*. [Online]. Available: http://newsfeed.time.com/2013/03/25/more-people-have-cell-phones-than-toilets-u-n-study-shows/. [Accessed: 12-May-2014].

[66] “Baby thinks print magazine is a broken iPad,” *Yahoo News*. [Online]. Available: http://news.yahoo.com/blogs/cutline/baby-thinks-print-magazine-broken-ipad-201148361.html. [Accessed: 12-May-2014].