Clarifying the Internet of Things by Defining the Internet of Devices

RICHARD L. RUTLEDGE, Georgia Institute of Technology
AARON K. MASSEY, Georgia Institute of Technology
ANNIE I. ANTÓN, Georgia Institute of Technology
PETER SWIRE, Georgia Institute of Technology

What observers have called the Internet of Things (IoT) presents privacy and security challenges for society. By surveying the literature on the IoT, we see that it rapidly evolved from industrial supply chain technologies to a merger of many technologies with distinct applications resulting in a single, difficult-to-define concept: the “thing.” What constitutes a “thing?” The term has referred to both sensed objects (things), such as the contents of a refrigerator, and to objects 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, are the objects that connect the world with the Internet. In this paper, we define both the IoT and the IoD and summarize the technologies from which they have evolved. In the process, we identify security and privacy challenges posed by these technologies. Technologists and policy makers must develop standards, protocols, identity management solutions, and governance to address these challenges. To this end, we develop a five-part general policy framework for detecting and responding to 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.

Categories and Subject Descriptors: A.1 Introductory and Survey, D.2 Software Engineering, D.2.1 Requirements Specifications

General Terms: Security, Standardization

Additional Key Words and Phrases: Internet of Things, Internet of Devices, privacy, security

1. 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 may not be Internet-connected and may not even be electronic equipment. They might simply be every-day objects represented as data. Initially, things were tagged with machine-readable identification technologies, like advanced Electronic Product Codes (EPC), Quick Response (QR) Codes, or Radio Frequency Identification (RFID) chips. However, IoT is now often used to refer to sensors or devices that directly connected to the Internet. Fekiet al. estimate that 50 to 100 billion devices will be connected to the Internet by 2020 [Feki et al. 2013]. This paper clarifies the definition of the Internet of Things and provides a consistent set of terms for technical elements of the IoT. In particular, we introduce a consistent vocabulary for technologists and policy makers seeking to mitigate security and privacy threats resulting from IoT technologies.

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:[1]

DOI:http://dx.doi.org/10.1145/0000000.0000000

> "The Internet of Things is a self-configuring and 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 is an idealistic, aspirational definition for the IoT. The IoT is not currently a fully self-configuring and adaptive system, nor are "all" things interconnected. 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. 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. A network of RFID readers can provide complete coverage of a facility. 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 [Singel 2004] such as skimming and eavesdropping by an adversary [Meingast et al. 2007].

Separating *things* from *devices* is critical for security and privacy analyses. 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" [Calo 2013]. 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. 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.

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.

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

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.

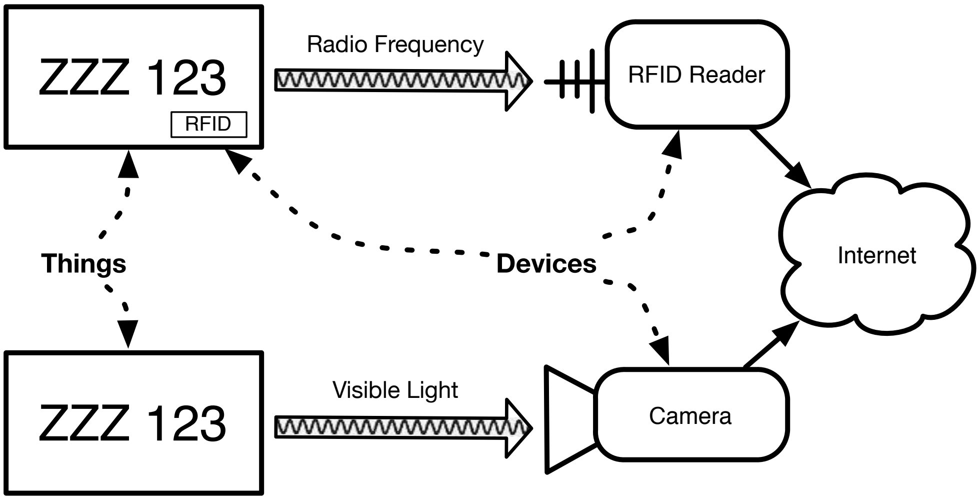

Figure 1: Differentiating Things and Devices in Systems for Tracking Cars by License Plates

We define *things* to be "any object about which a device collects data or upon which a device operates." We define *devices* as "the technologies that collect data from their environment, interact with their environment, and communicate directly with a network." We define the *Internet of Devices (IoD)* to be "the collection of devices that are capable of operating, either directly or indirectly, over the Internet." Colloquial usage of the word 'thing' does not imply network communications. Smarter devices capable of interacting with their environment are now commonly considered to be a part of the IoT. Consider thermostats that 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. In 2014, Google purchased Nest, a company that makes such a device, for $3.2 billion dollars [Wohlsen 2014]. An author for Wired described the purchase as a marker that the IoT has become mainstream [Wohlsen 2014]. Some Nest users also use Android phones, also produced by Google. When those users leave work 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.

Along with the need for conceptual clarity by separating things from devices, 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.** Identification-only devices that are physically attached to things (e.g. RFID)
2) **Remote sensors.** Devices that can recognize and identify things remotely (e.g. cameras with product recognition software)
3) **Smart devices.** Devices with 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.** Devices that are functionally similar or equivalent to the desktops, laptops, and servers we use today.

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. Finally, we discuss briefly the policy implications of different kinds of data generated or transmitted by devices on the IoD.

The remainder of this paper is organized as follows. Section 2 introduces our terminology for the constituent elements of the IoD. In Section 3, we present our survey methodology. In Section 4, we discuss the results of our survey. In Section 5, we present our framework for examining the key security and privacy challenges of the IoD. We differentiate IoD privacy and security concerns from other privacy and security concerns in Section 6. Finally, we summarize our analysis in Section 7.

2. TERMINOLOGY

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. Finally, devices are designed, final products. If a consumer attaches a generic identification device to something, it remains a thing attached to a device rather than becoming a device.

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" [Privat 2012], which we discuss in more detail in Section 4. 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, it becomes a thing in the IoD. Moreover, if two or more devices are used to collect the same data about a single object, 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. Components include but are not limited to sensors, articulators, and identifiers.

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 its network. Example sensors include temperature and location sensors.

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. Example devices employing articulators include automated door locks and smart grid power switches. A less obvious example of a device with an 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.

IoD Communications Protocol: A system of rules for data exchange across a network and between devices. Some devices may support multiple, simultaneous communications protocols over multiple 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 and that can either effect or monitor the state of physical, 'real-world' objects. IoD networks do not have to use communications protocols common to the Internet. Instead, they may choose to use a proprietary protocol for communications.

Internet of Devices: A network of interconnected IoD Networks.

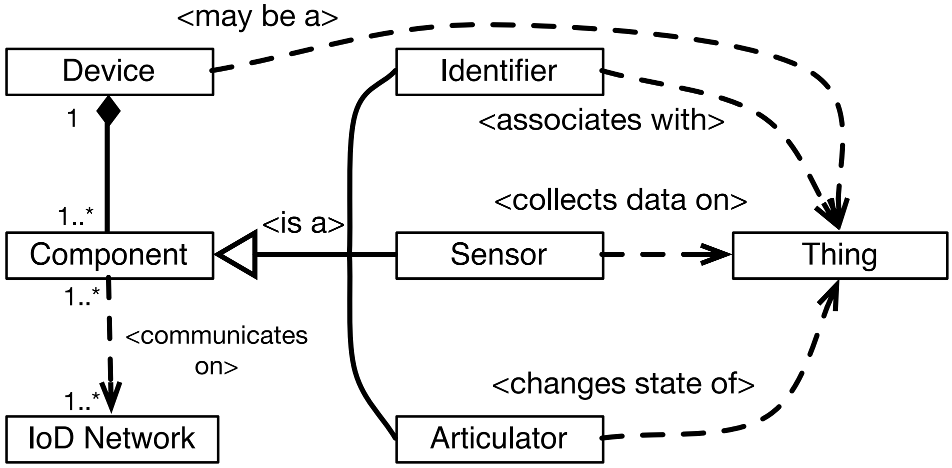

Figure 2: Relationship between Devices, Components, and Things

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 represent 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 [Kerr 2013]. Human-implantable RFID chips wirelessly and automatically identify people for a variety of purposes, including medical records and payment systems. Some people like the convenience of being easily identifiable to computers.

People might not even be aware of 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. The owner may now be considered a thing when they wear the coat with the RFID tag still attached. 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 allowing them to be categorizing as "things" in our definition. Instead, we seek to accurately describe the relationship between IoD devices and the people who 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 refrigerator examples illustrated in Figure 3. Example A contains no devices. Both the refrigerator and the milk carton are simple things. Example B shows a refrigerator with a simple camera installed, which records the refrigerator's contents. This makes the refrigerator a device, but the milk carton remains a thing. Example C shows a refrigerator with an RFID reader installed, which once again makes the refrigerator a device. This example also shows a milk carton with an RFID tag. If this tag was part of the supply chain management of the grocery store, then the milk carton is a device. However, if this RFID tag was simply taped to the milk carton by the owner rather than being built into the milk carton, then the milk carton remains a thing and the RFID tag is a standalone device.[2] In Example D, the refrigerator is a device by virtue of both an articulator and a sensor. Note that the smart refrigerator does not have to directly connect to the Internet so long as a process exists for transmitting information from the device, such as a record of food items stored in it or a record of cooling efficiency, to the Internet. For example, a company may produce a smart refrigerator that exports data as plain text files that can be uploaded by the owner and examined by tech support. Alternatively, the smart refrigerator may sync with the Internet through an application on a mobile device such as a cell phone.

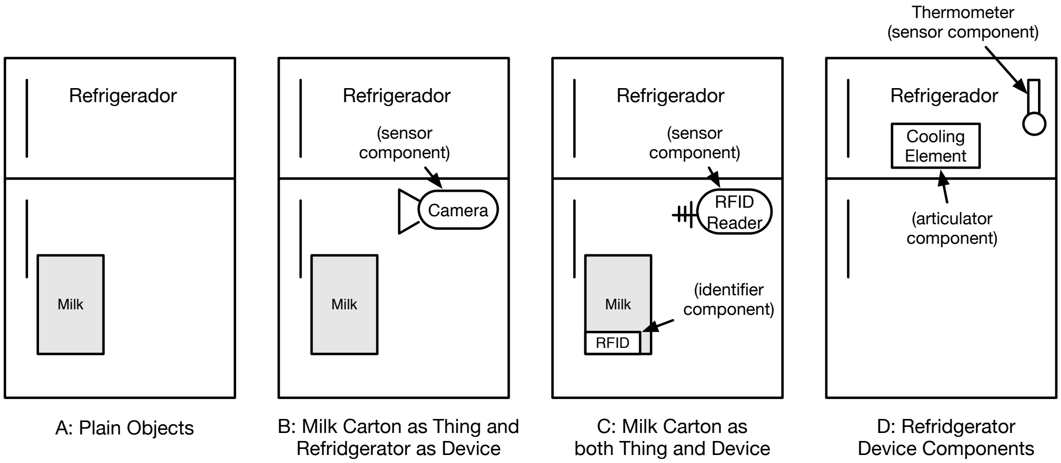

Figure 3: Smart Home Technology Examples

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 primary 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. Clear

[2] Consider the milk carton in B. If it was printed with a QR code, then it would become a device rather than a thing.

definitions can reduce confusion between engineers (focused on devices) and policy-makers (focused on things).

3. METHODOLOGY

To maximize the number of primary sources found while minimizing induced bias, we elected to perform a systematic literature review (SLR) [Keele 2007]. We performed an automated search of research repositories with documented search terms for repeatability and auditability. We pruned this returned set by manual inspection for papers applicable to our research question and augmented with some targeted manual searches [Kitchenham et al. 2009]. Further research has supported the reliability of SLRs [MacDonell et al. 2010] and underscored the importance of documenting the search process [Kitchenham et al. 2011].

The IEEEXplore database, the ACM Digital Library database, and the Google Scholar database were each searched with the following keywords: "Internet Things Privacy", "IoT Privacy", "Internet Things Security", and "IoT Security". The security keyword was added to increase the scope of the search results, but results that did not also consider privacy were discarded. The titles and abstracts from the top 50 hits from each repository were manually inspected and relevant articles were downloaded for further inspection. The initial search yielded 22 articles. Each article was examined, and any relevant works cited by those papers were also considered for inclusion in our survey.

Subsequent to the initial search, the Social Science Research Network (SSRN) and JSTOR repositories were identified as sources for policy and legal articles. The same search procedures were applied to these repositories, with poor results. Title and abstract inspection excluded a high percentage of the resulting articles as irrelevant to our topic or of poor quality. Moreover, many known relevant policy and legal papers were not returned in the results. We consulted with a co-author and legal expert on applying Kitchenham's survey methodology to legal and policy repositories to understand the root causes for the noisy resultant set. After reviewing our results and querying the LexisNexis legal database, we identified the following attributes that make legal and policy repositories resistant to systematic literature review.

Unlike computer science repositories, legal and policy repositories are not keyword indexed. Their reliance on content indexing suffices when the search terms are highly distinguished, such as 'civil liberties'. However, adding common words (sometimes called stop words) to the search terms may not significantly alter the results. And the IoT is a highly distinguished conceptual model identified by a phrase of effectively ubiquitous words: 'Internet' and 'Things'. The result of our search queries was very similar to searches on 'Privacy' alone. In contrast, a computer science article concerning the IoT would include both 'Internet' and 'Things' as keywords; resulting in the article's elevation in the ranking of search results.

Another essential difference between computer science and legal and policy repositories is that a relatively small number of computer science databases can be considered to be comprehensive to the field. The legal and policy fields' diversity of publication venues complicate the identification of a comprehensive set of relevant papers. If a systematic literature review is conducted on a non-comprehensive repository set, then important perspectives on the research question may be overlooked. Moreover, the requirement to apply the same search to each repository assumes a significant degree of symmetry in search capability and algorithms. Although this assumption may hold across technical domains, less technological domains may lack established procedures, such as use of keywords, that aid automated indexing.

4. SURVEY RESULTS

As computing devices became smaller, more power efficient, and cheaper to produce, they transformed the Internet of Tagged-Things into the Internet of Smart-Things, with a corresponding increase in security and privacy risks. We present our survey results using a categorization based in part on this transition. We begin with devices used for inventory and supply chain management. These devices are identification-only and focus on radio frequency identification (RFID) devices, where the term “Internet of Things” originated. Second, we discuss remote sensors focusing on wireless sensor networks (WSNs), which mark the transition from devices that only provide identity information to devices that report on and interact with their environment. Third, we discuss consumer devices in the smart home and smart office contexts. Fourth, we examine wearable and ubiquitous computing. In this case, devices are essentially application-specific computers. 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 2. 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.

4.1 Inventory and Supply Chain Management

Kevin Ashton coined the phrase “Internet of Things” in a 1998 presentation to Procter and Gamble [Santucci 2009]. 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 [Santucci 2009].” 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 [Weber 2010]. Servers maintain information about tagged things and index them on the unique RFID identifier, which allows readers to query servers for more information [Weber 2010].

Ashton and others founded the MIT Auto-ID Center that envisioned objects (things) tagged, identified, and tracked by RFID [Feng and Fu 2010]. 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 [Weber 2009]. 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.

IoT technologies and standards, such as EPC global, were initially based on Internet standards. Consider object naming and identification. EPC global has an

Object Naming Service (ONS), which is based on the Internet's Domain Name Service (DNS) [Weber 2009]. 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 [Weber 2009]; he believes similar concerns will also apply to EPC global. For example, VeriSign currently operates the ONS directory service root node for EPC global and, as a result, has a great deal of practical influence over how EPC global 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[Weber 2009]."

At least two IoT (IoD) research groups proposed adapting the Platform for Privacy Preferences (P3P), an Internet standard for expressing privacy preferences [Tao and Peiran 2010; Ukil et al. 2012]. 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) [Tao and Peiran 2010]. 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 [Tao and Peiran 2010]. Ukil et al. also identified the individual data producer and data consumer as stakeholders in their negotiation-based privacy preservation technique [Ukil et al. 2012]. 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 [Ukil et al. 2012]. The Negotiation Module is governed by privacy policies that are, in turn, based upon privacy law [Ukil et al. 2012]. 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 [Yu et al. 2004; Hogben 2002; Karjoth et al. 2003; Reay et al. 2007; Cranor 2012]. In fact, some proposed P3P adaptations, such as the automated negotiation modules proposed by Ukil et al., 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 [Cranor 2012].

Machara et al. propose to insert a Context Manager Middleware layer into the IoT (IoD) [Machara et al. 2013]. Rather than starting from a P3P baseline, the authors develop a context-oriented model of an agreement between producers and 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, such as the data to be read by the consumer [Machara et al. 2013]. If a match cannot be made, then the data is not made available to the would-be consumer [Machara et al. 2013]. One of the advantages of this approach over P3P adaptations is the ability to handle dynamic modifications to the producer and consumer contracts. Although these meta-models have been validated with the Eclipse Modeling Framework (a tool for checking model consistency) their complexity may be a significant obstacle to broad adoption.

The IoT inherits existing fundamental security concerns from the Internet. These concerns are exacerbated by the greatly expanded scope and scale of the IoT. An adversary may have little incentive to track an RFID encoded milk carton (thing), but an RFID encoded wallet (thing) linked to an individual consumer may prove more valuable. Zhu et al. considered the security of connections between RFID tags (identifiers), readers (sensors), and backend systems such as the Object Name Service (ONS) [Zhu et al. 2012]. 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 EPC global RFID system exposed to the public domain [Khoo 2011]. 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 [Khoo 2011]:

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 [Khoo 2011]." 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).

4.2 Wireless Sensor Networks

Wireless sensor networks (WSNs) are networks of sensor devices that connect with each other, and possibly the Internet, wirelessly. WSNs may have begun the confusion with the terms "thing" and "device." The word "thing" applied initially to RFID chips, where the inventory item and the unique identifier were physically combined. Devices and things are entirely different objects in WSNs. This difference in terminology is critical for policies governing the IoD. The networking of sensors should mean that policy analysis on the sensors themselves, which are devices, rather than on the particular things being sensed.

Sensor networks have a long history. 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 [Chong and Kumar 2003]. Additional DARPA programs such as Sensor Information Technology further developed robust, ad hoc networking and distributed information processing [Chong and Kumar 2003]. At the same time, advances in micro-electromechanical systems decreased the size, power consumption, and cost of sensors (devices) while simultaneously increasing their range. The addition of wireless communications technology fundamentally transformed sensor networks and enabled the transition from DSNs to WSNs.

A typical WSN is composed of a large number of self-contained, communicating sensor packages (devices) [Akyildiz et al. 2002]. 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 [Akyildiz et al. 2002]. Each individual sensor node (device) may have minimal computational resources, but the aggregate network may have considerable computational capability [Culler et al. 2004]. Applications for WSNs include military, security, and environmental monitoring. Yick et al. categorize WSN applications as either Tracking or Monitoring [Yick et al. 2008]. 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 [Culler et al. 2004]. 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 [Culler et al. 2004]. These communications strategies provide resiliency in the complete system given individual sensor node failures and communications interference from physical obstacles. From a policy perspective, this resiliency means that a single node may be communicating data it did not generate and describing things it cannot sense directly.

The European Union's IoT Architecture (IoT-A) project proposed a reference architecture that provides interoperability between RFID systems and WSNs. In this context, Gessner et al. consider the requirements for object resolution functions [Gessner et al. 2012]. 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 [Gessner et al. 2012]. 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 [Gessner et al. 2012]. The authors do not specify this module in 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. The first four modules may enable effective policies to govern information sharing in these networks, but the trust module is both critical to this effort and incomplete. In addition, RFID tags lack the computational resources to meaningfully participate in elements of the proposal, such as Authentication.

Privat is one early researcher who wrote about the dangers of conflating "things" and "devices." His proposal to include mundane, non-communicating objects as "things" in the Internet of Things was based on the remote tracking capabilities of sensors [Privat 2012]. He termed such objects sense-able things. The implications of this difference are critical for policy analysts and should not be overlooked by engineers. We now consider a standard IoT transaction and compare it to a hypothetical extended IoT transaction.

Comparing two approaches to smart refrigerator technologies illustrates the subtle effects on privacy and security within a single technical domain. In both examples, the smart refrigerator detects that the refrigerator has run out of milk and orders more. In the first scenario, the smart refrigerator is equipped with an RFID sensor and the milk carton has an embedded RFID chip. The consumer opens the refrigerator, removes a milk carton, and empties it into a glass. He then discards the empty milk 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 more from the Acme Corporation for delivery in the morning.

In the second scenario, the smart refrigerator is equipped with multiple interior cameras. The consumer opens the refrigerator, removes a milk carton, empties it into a glass, and discards the empty carton into a waste bin. Overnight, the refrigerator uses cameras to examine its contents. Its image detection and recognition software is unable to match with a milk container. It places an order for more from the Acme Corporation for delivery in the morning. The carton in the first refrigerator is both a thing and a device because data about it is collected and because it has an embedded electronic communications device, an RFID tag. The carton in the second refrigerator is also a thing, but it is not a device because has never been designed to include any components. Yet, the sensor in the second refrigerator detects and recognizes the milk carton. From a policy standpoint, the effect is the same as the one in which the milk carton was explicitly designed as a device.

Researchers who define “things” on the IoT as uniquely identifiable may not consider the milk carton in the second example to be a “thing” because it is not uniquely distinguishable. In the first example, the RFID chip would allow the smart refrigerator to distinguish between two milk cartons produced by the same company with the same external appearance. The refrigerator in the second example would not be able to do this.

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 things under this paper’s terminology Neither of these examples suffers from the quasi-identification of the refrigerator example. In each case, the object is uniquely identified.

Another concern with WSNs is that a security-compromised node (device) can potentially reveal information from the entire system. Data on WSNs is often 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) [Oleshchuk 2009]. SMCs allow collaborative computation without any party divulging its own input. SMCs can solve problems like Yao’s Millionaire problem: How can two millionaires determine which is the richest without revealing their own net worth? In IoD terms, the problem is to determine 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 [Oleshchuk 2009]. The secure sum protocol would allow a set of sensors to compute the sum of their values without disclosing any of their individual values. The last SMC algorithm 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.

4.3 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,[3] 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) [Metz 2013]. 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,[4] a location-based social network [Metz 2013]. 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 [Anon n.d.]. This project's goal is to enable research into how a controlled home environment can improve health, well-being, 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 [Cook et al. 2003]. 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 currently 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 [Chan et al. 2008].

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 accurately assess which appliances a resident might be using and when those appliances are used. This assessment is based

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

entirely on the timing, amount, and signature of the power events in the house. Similar to WSNs, these assessments allow tracking of "things" that were not originally designed with the intention of assisting in their own tracking. Interoperability between appliances and a smart grid will also allow washing machines and dishwashers to automatically start when electricity prices are lowest [Kominers 2012b]. However, such interoperability must be carefully designed to mitigate 'Perverse Results' [Kominers 2012a] in which local optimizations reduce global efficiency. If every dishwasher started at the same cost threshold, then the increased demand would cause the price to rise.

The policy implications of smart homes are profound. Homes are intimate spaces that have traditionally received legal protections in many jurisdictions 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 collected and stored inside the home.

Kanuparthi et al. addressed some privacy and security threats from the smart home with Physical Unclonable Functions (PUFs) [Kanuparthi et al. 2013]. PUFs are the hardware equivalent of a cryptographic one-way function and can be used in a challenge-response protocol [Kanuparthi et al. 2013]. 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. For example, if PUFs were applied to consumer electronics, correct assessment of electricity usage may be impossible. 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" [Kanuparthi et al. 2013]. 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 ensure privacy and security in this environment, Kanuparthi proposes to integrate PUFs into IoT sensors and use PUFs for device identity management [Kanuparthi et al. 2013]. 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 [Kanuparthi et al. 2013]. 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 [Mayer et al. 2012]. 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 [Mayer et al. 2012]. A visualization aide of this type may have lead to an earlier detection of an LG smart television leaking personal viewer data [Kelion 2013b]. 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.

4.4 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 [Federal Trade Commission n.d.]. 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 (device), an optical head-mounted display (articulator) with a camera (sensor) and microphone (sensor), might 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 future, we are no closer to understanding the impact they will have on individual privacy and security.

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 geo-locational devices enables portable access to location information. These devices include GPS satellite tracking, cellular tower triangulation, and Wi-Fi finger printing 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 [Elkhodr et al. 2012]. 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 [Elkhodr et al. 2012]. 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 [Elkhodr et al. 2012]. 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 [Elkhodr et al. 2012]. Fifteen of these 30 applications were sending geographic location information to remote advertising servers. Seven of these applications even provided the phone's unique hardware identifier [Elkhodr et al. 2012], which would allow data from one application to be linked to data from another application that also has 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 [Liu et al. 2012]. The phone's LBS request services through the middle-ware that relays the request to the service provider through a pseudonymous account [Liu et al. 2012]. Hu et al. also propose a middle-ware layer of software to provide emergency access to LBS data [Hu et al. 2011], 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 [Liu et al. 2012], 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 Wi-Fi and cellular communications networks can act as hubs and allow other devices that do not have Wi-Fi 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 [Lo et al. 2006]. PANs offer a natural architecture for technical measures to enforce policy goals, such as protecting privacy and security.

4.5 An Evolving Internet

A final 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 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 [Wang and Wen 2011]. 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 [Wang and Wen 2011]. 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 [Wang and Wen 2011]. 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 [Huang et al. 2012]. 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 [Huang et al. 2012]. For data access, the authors describe a context aware, k-anonymity [Sweeney 2002] policy and filter. They illustrate this with an example of a producer/consumer pair who are colleagues and the data 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 [Huang et al. 2012]. 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 Section 4.2, will not scale into the IoT (IoD) since these techniques require the naming of individuals to be granted or denied access [Evans and Eyers 2012]. They maintain that ensuring consistent implementation of discretionary access would be impractical in the highly dynamic environment of the IoT. They propose to use techniques from Information Flow Control (IFC) to label data packets directly 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 [Evans and Eyers 2012]. 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 requirement for a TCB remains necessary. Also, a comprehensive ontology for tagging privacy-related data would be difficult to achieve in advance and would need to be maintained indefinitely.

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 [Weber 2010]. He rejects national regulation as not meeting the IoT globalization requirements [Weber 2010]. He acknowledges that neither international agreement nor self-regulation alone would be practical to implement and acceptable to preserving privacy [Weber 2010]. He recommends a form of "co-regulation" in which government sets a general framework elaborated by the private sector [Weber 2010]. Weber also notes the special difficulties in achieving globalization given the differing notions of privacy in various regions of the world [Weber 2010].

5. A FRAMEWORK FOR EVALUATING SECURITY AND PRIVACY IN THE IOD

The Internet of Devices presents security and privacy challenges for software engineers, regulators, and policy makers. In this section, we describe a framework based on our survey of existing research. We begin with an examination of inputs and outputs as a starting point for driving policy decisions for all devices on the IoD in Section 5.1. We then consider in Section 5.2 the identifiability and linkability of the information transmitted. Not all information poses the same risk for security or privacy violations. In Section 5.3, we examine devices across the entire IoD spectrum, and we consider some of the implications for future IoD development.

5.1 Towards Heuristics based on Inputs and Outputs

Perhaps the simplest heuristic for examining security and privacy on the Internet of Devices is a simple two-by-two matrix as shown in Figure 4. Devices that accept inputs may present security concerns for the device. Each input accepted, if mishandled by the device, could result in the device being compromised. For example, flooding a wireless door lock with noise could jam it and prevent its intended operation. Devices that produce output may present privacy concerns. Each output produced could include information about a person that might compromise their privacy. For example, a door lock that outputs a log of entries and exits could provide a wealth of information on the homeowner's comings and goings.

Applying this heuristic yields four types of devices. Type 1 devices may present 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 present privacy concerns. Type 3 devices accept inputs, but produce no outputs. Type 3 devices may present security concerns, but they cannot present privacy concerns. Type 4 devices accept no inputs or outputs and, as a result, present no security or privacy concerns.

This heuristic is easy to interpret, but is it useful for evaluating security and privacy in IoD devices? Does each type represent a set of realistic devices? 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 or shoes. These tools are so simple that they are considered to be everyday objects.

Type 2 and Type 3 devices are more interesting. A Type 2 device that accepts no inputs but produces outputs may correspond to a sensor as defined in Section 2. Similarly, a Type 3 device that 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 single output, the device would need to be considered a Type 1 device. Similarly, if a Type 3 device accepted even a single input, it would 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. Many 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	$S \land P$ 1	$\neg S \land P$ 2
No Output	$S \land \neg P$ 3	$\neg S \land \neg P$ 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 special cases. Most types of devices run standard communications protocols that are inherently bidirectional. In order to both accept inputs and produce outputs, devices must, by definition, 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 poses no threat to privacy so long as the data remains in the device. 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.

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?

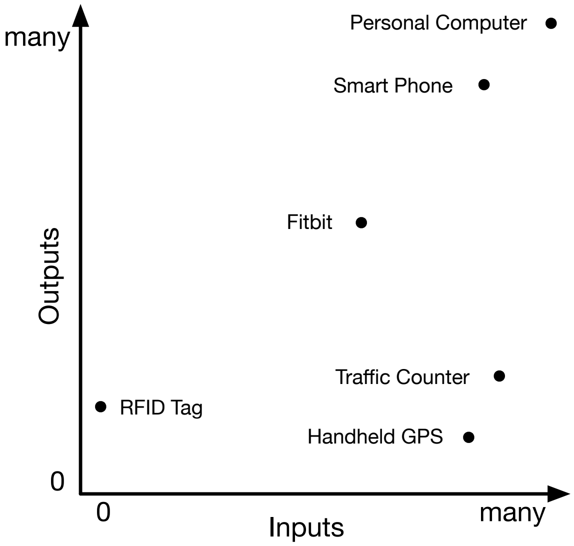

Figure 5: A continuous model for security and privacy concerns for devices on the Internet of Devices

Extremely challenging analysis scenarios are easy to construct. Devices may accept many sensitive inputs, produce many sensitive outputs, and communicate broadly with people or other devices. They might communicate autonomously or with limited human input; for example, self-driving cars, autonomous drones, or citywide self-regulating traffic systems. But for all of this complexity and all of this communication, a simple examination of inputs and outputs remains a useful starting point – heuristically focusing attention on where security and/or privacy issues call for further attention.

5.2 Towards Heuristics for Identifiability and Linkability

Information varies in its sensitivity for privacy and security. For example, personally identifiable information (PII) is more relevant to privacy than non-identifiable information, and information that enables root access is especially relevant in security. The value of data collected or processed by IoD devices spans the entire spectrum from innocuous to extremely sensitive. The sensitivity of this data should inform the security and privacy analysis of the device. For example, data that explicitly identifies a particular individual is more sensitive than data that must be linked to an external data source to identify that same individual.

In this section, we examine the role of identifiability and linkability in a security and privacy analysis of IoD devices. Inputs and outputs are, of course, not the only factors that need to be examined. We have at least three things that may need to be differentiated: (1) devices that communicate with people and devices that communicate with other devices; (2) individuals and groups; and (3) automated communications and mediated communications. 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. Individuals and groups often have different privacy protections. Access of a particular individual's information may be justified when access of an entire group's information is not. Devices that cannot communicate without human interaction (i.e. mediated devices) 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. Finally, we may need to consider the information itself as a potential source for privacy and security concern. Some information, such as a medical history, directly reveals details about a person's life. Other information, such as a social security number (SSN), enables linking together data collected separately.

Mere complexity does not imply greater privacy or security risk. As an initial heuristic, the sensitivity of data collected or processed by IoD devices varies based on whether it (1) identifies one or more individuals or (2) does not identify any individual. Some of the simplest IoD devices are merely identifiers. RFID-enabled toll collection devices, for example, make it easy for motorists to travel without long delays on toll roads. These devices identify the user to be billed automatically upon traveling through the tollbooth. That identification can be sensitive information, such as when it has been used in criminal and civil court proceedings to establish the location of individuals. For comparison, consider a device that measures moisture for houseplants. This device is more complex than the RFID chip, because it must also accept some input, process that input, and produce an appropriate output. Despite additional complexity, the data involved poses less of a security or privacy risk, because no individual is identified.

Another useful extension to our framework for privacy analysis considers the domain of sensor data collected. The collection domain consists of the subject of the data collection together with the context in which the data is collected. First, we will consider the sensor subject. Sensor data has the potential to compromise the secrets of its subject. When that subject is an individual person, we consider the compromised secret to be a potential privacy harm. Thus facial recognition software in a public plaza is a threat to privacy and the storage and access to this data must be analyzed further in a big data context for privacy threat mitigation. When the sensor subject is an entity, such as an organization, compromised secrets do not directly harm privacy. Organizations do not have privacy. The loss of secrets in a government organization can result in a threat to national security. Similarly, such loss in a corporate organization is a potential source in the leak of trade secrets. However, organizations have an explicit association with individual people as members, employees, etc. Here the context in which the data is collected helps to determine the risk of inferring data regarding associated people.

Although inanimate objects cannot possess secrets and suffer privacy harms, they can be a risk to secrets as the sensor subject through either explicit or implicit association with people. Consider a smart meter monitoring the power utilization at a factory. If the factory is related to a corporation through ownership, then the sensor has the potential to be a threat to company secrets, but has no privacy implications. However, the context of the data sensor is essential for a sound analysis of privacy risks. The association of objects to people can be a dynamic relationship. For example, an RFID chip embedded in a consumer good is associated through ownership with a variety of corporate entities throughout the manufacturing and supply chain management process. Throughout this stage, the RFID represents a potential threat to company trade secrets. But when a consumer purchases the product, the relationship changes to a personal one, and the potential threat shifts to a risk to privacy.

The common use of standard bidirectional communications protocols means that many IoD devices will not be classifiable as input-only or output-only. 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 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. The device warrants a privacy analysis due to the media titles and an implicit association of the device with a user account, but the volume of data flow by itself does not indicate the level of risk.

5.3 Device Categories, From Simple to Complex

Our framework, which focuses on inputs and outputs, works especially 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.

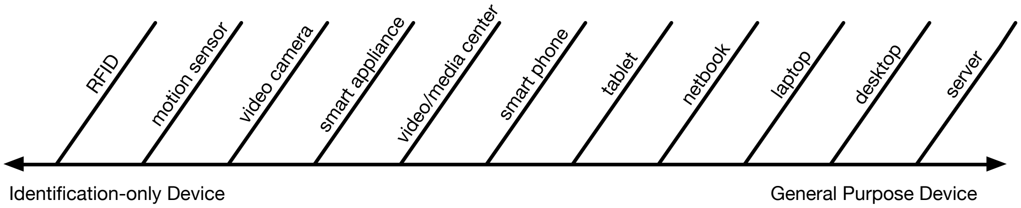

Figure 6: Spectrum of Devices in the IoD

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 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 [Perry n.d.; Rolfe n.d.; King n.d.] 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 the driver or someone walking a 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. Complete analysis is beyond the scope of this paper and involves general privacy and security issues that are not specific to the IoD. 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, with their attendant privacy and security risks.

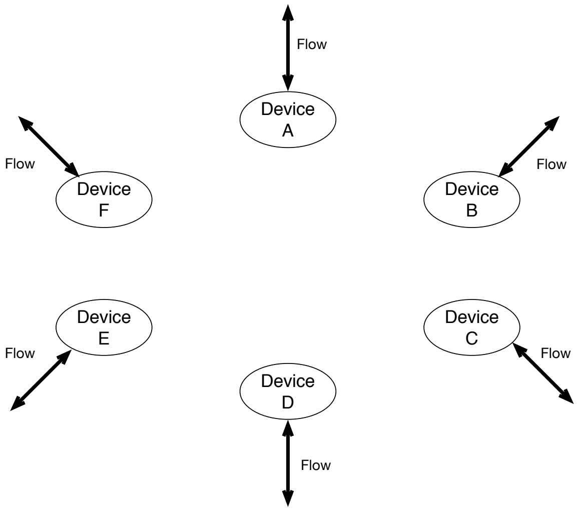

Figure 7: Stand-alone Devices

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 [Kelion 2013a]. 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 was attached to the TV, it sent 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) often engages only 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 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 [Samsung n.d.]." A flow diagram would look more like Figure 8 and the number of flows to be analyzed increases with the square of the number of devices.

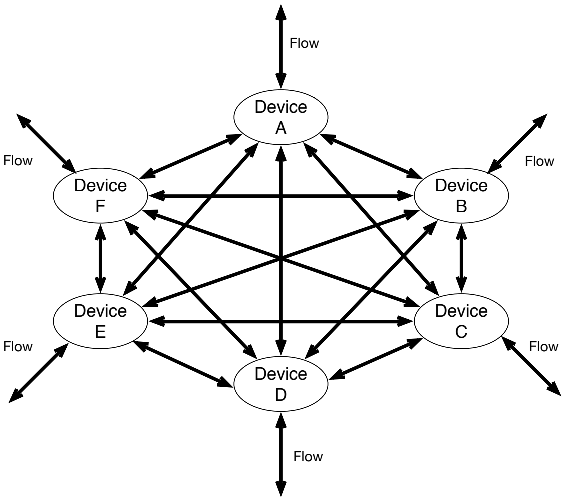

Figure 8: Collaborating Devices

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 enmeshed system rather than just the single device.

6. DIFFERENTIATING PRIVACY CONCERNS

The Internet of Devices is not the only challenge for privacy and security. It is important, therefore, to differentiate between the challenges posed by the IoD and other extant challenges, such as Big Data, Cloud Computing and Robotics. 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 [Mayer-Schönberger and Cukier 2013]. 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 [Qin et al. 2013]. Autonomous devices, whether simple or complex, could justifiably be thought of as robots or drones, which also have a separate scholarship related to privacy [Calo 2010]. Similarly, ubiquitous computing is another field in which scholars examine security and privacy concerns [Camp and Connelly 2007]. 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 [Mayer-Schönberger and Cukier 2013]. 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" [Mayer-Schönberger and Cukier 2013]. Often, the ability to do things at a large scale will depend entirely upon data collected using IoD devices. Consider a large retailer, like Wal-Mart, 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 [Mayer-Schönberger and Cukier 2013]. 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" [Ohm 2010]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 simple IoD 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 [Qin et al. 2013]. 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 the residents use, 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 [Greveler et al. 2012]. Third party doctrine is a particular concern for privacy in cloud computing services [Harper 2008], [Soghoian 2009]. Cloud computing is also further exacerbated by location-based jurisdiction issues for legal systems all over the world [Desai 2013]. 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 [Schneier 2012]. 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" [Calo 2011]. 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 [Balch and Arkin 1994]. 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 [Camp and Connelly 2007]. 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 [Camp and Connelly 2007; Price et al. 2005; Strahilevitz 2008]. The solutions and mitigations for those challenges may apply to IoD challenges as well.

7. SUMMARY

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 an easier examination of security and privacy concerns. The more complex the device, the more complex the potential security and privacy concerns may be, and the greater the interactions with existing literatures including Big Data and Cloud Computing.

We also provide a simple framework for analyzing security and privacy concerns for IoD devices. 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 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 Telecommunications Union claims that at some point in 2014 cell phones will outnumber people [Anon n.d.]. The United Nations claims that more people have access to cell phones than toilets [Wang n.d.]. 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 [Anon n.d.]. 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

AKYILDIZ, I.F., SU, W., SANKARASUBRAMANIAM, Y. AND CAYIRCI, E., 2002. Wireless Sensor Networks: A Survey. *Computer networks*, 38(4), pp.393–422.

ANON, 2014: Mobiles "to outnumber people." *BBC News*.

ANON, About AHRI.

ANON, Baby thinks print magazine is a broken iPad. *Yahoo News*.

ANON, 2008. Commission Staff Working Document, Future Networks and the Internet – Early Challenges regarding the "'Internet of Things.'" *Samsung Electronics America*.

BALCH, T. AND ARKIN, R., 1994. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1), pp.27–52.

CALO, M.R., 2013. Digital Market Manipulation. *University of Washington School of Law Research Paper*, (2013-27).

CALO, R., 2010. People Can Be So Fake: A New Dimension to Privacy and Technology Scholarship. *Penn State Law Review*, 114(3).

CALO, R., 2011. Robot Ethics: The Ethical and Social Implications of Robotics (Intelligent Robotics and Autonomous Agents series). In P. Lin, G. Bekey, & K. Abney, eds. The MIT Press.

CAMP, L.J. AND CONNELLY, K., 2007. Digital Privacy: Theory, Technologies and Practices. In A. Acquisti, S. D. C. di Vimercati, S. Gritzalis, & C. Lambrinoudakis, eds. Taylor & Frances, New York, NY.

CHAN, M., ESTÉVE, D., ESCRIBA, C. AND CAMPO, E., 2008. A Review of Smart Homes-Present State and Future Challenges. *Computer Methods and Programs in Biomedicine*, 91(1), pp.55–81.

CHONG, C.-Y. AND KUMAR, S.P., 2003. Sensor Networks: Evolution, Opportunities, and Challenges. *Proceedings of the IEEE*, 91(8), pp.1247–1256.

COOK, D.J. ET AL., 2003. Mavhome: An Agent-Based Smart Home. In *Pervasive Computing and Communications, 2003. (PerCom 2003). Proceedings of the First IEEE International Conference on*. pp. 521–524.

CRANOR, L.F., 2012. Necessary but Not Sufficient: Standardized Mechanisms for Privacy Notice and Choice. *J. on Telecomm. & High Tech. L.*, 10, p.273.

CULLER, D., ESTRIN, D. AND SRIVASTAVA, M., 2004. Guest Editors' Introduction: Overview of Sensor Networks. *Computer*, 37(8), pp.41–49.

DESAI, D.R., 2013. Beyond Location: Data Security in the 21st Century. *Communications of the ACM*, 56.

ELKHODR, M., SHAHRESTANI, S. AND CHEUNG, H., 2012. A Review of Mobile Location Privacy in the

Internet of Things. In *ICT and Knowledge Engineering (ICT Knowledge Engineering), 2012 10th International Conference on*. pp. 266–272.

EVANS, D. AND EYERS, D.M., 2012. Efficient Data Tagging for Managing Privacy in the Internet of Things. In *Green Computing and Communications (GreenCom), 2012 IEEE International Conference on*. pp. 244–248.

FEDERAL TRADE COMMISSION, Internet of Things - Privacy and Security in a Connected World.

FEKI, M.A., KAWSAR, F., BOUSSARD, M. AND TRAPPENIERS, L., 2013. The Internet of Things: The Next Technological Revolution. *Computer*, 46(2), pp.24–25.

FENG, H. AND FU, W., 2010. Study of Recent Development about Privacy and Security of the Internet of Things. In *Web Information Systems and Mining (WISM), 2010 International Conference on*. pp. 91–95.

GESSNER, D., OLIVEREAU, A., SEGURA, A.S. AND SERBANATI, A., 2012. 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*. pp. 998–1003.

GREVELER, U., JUSTUS, B. AND LOEHR, D., 2012. Multimedia Content Identification Through Smart Meter Power Usage Profiles. *Computers, Privacy and Data Protection*.

HARPER, J., 2008. Reforming Fourth Amendment Privacy Doctrine. *American University Law Review*, 57, p.1381.

HILDEBRANDT, M., 2009. *Profiling and the Rule of Law*, Rochester, NY: Social Science Research Network.

HOGBEN, G., 2002. A Technical Analysis of Problems with P3P 1.0 and Possible Solutions. In *Position paper, W3C Workshop on the Future of P3P*.

HU, C., ZHANG, J. AND WEN, Q., 2011. 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*. pp. 192–195.

HUANG, X., FU, R., CHEN, B., ZHANG, T. AND ROSCOE, A.W., 2012. User Interactive Internet of Things Privacy Preserved Access Control. In *Internet Technology And Secured Transactions, 2012 International Conference for*. pp. 597–602.

IGLEZAKIS, I., 2013. *Regulation Models Addressing Data Protection Issues in the EU Concerning RFID Technology*, Rochester, NY: Social Science Research Network.

KANUPARTHI, A., KARRI, R. AND ADDEPALLI, S., 2013. 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: ACM, pp. 61–64.

KARJOTH, G., SCHUNTER, M., VAN HERREWEGHEN, E. AND WAIDNER, M., 2003. Amending P3P for Clearer Privacy Promises. In Database and Expert Systems Applications, 2003. Proceedings. 14th International Workshop on. pp. 445–449.

KEELE, S., 2007. *Guidelines for Performing Systematic Literature Reviews in Software Engineering*, Technical report, EBSE Technical Report EBSE-2007-01.

KELION, L., 2013a. LG investigates Smart TV `unauthorized spying' claim. *BBC News*.

KELION, L., 2013b. Lg Investigates "Spying" Smart Tvs. *BBC News*.

KERR, I.R., 2013. *The Internet of People? Reflections on the Future Regulation of Human-Implantable Radio Frequency Identification*, Rochester, NY: Social Science Research Network.

KHOO, B., 2011. 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*. pp. 709–712.

KING, R., U.S. testing crowd-scanning facial recognition system.

KITCHENHAM, B. ET AL., 2009. The Impact of Limited Search Procedures for Systematic Literature Reviews #x2014; a Participant-Observer Case Study. In *Empirical Software Engineering and Measurement, 2009. ESEM 2009. 3rd International Symposium on*. pp. 336–345.

KITCHENHAM, B., BRERETON, P., LI, Z., BUDGEN, D. AND BURN, A., 2011. Repeatability of Systematic Literature Reviews. In *Evaluation Assessment in Software Engineering (EASE 2011), 15th Annual Conference on*. pp. 46–55.

KOMINERS, P., 2012a. *Interoperability Case Study: Internet of Things (IoT)*, Rochester, NY: Social Science Research Network.

KOMINERS, P., 2012b. *Interoperability Case Study: The Smart Grid*, Rochester, NY: Social Science Research Network.

LIU, J., HU, X., WEI, Z., JIA, D. AND SONG, C., 2012. Location Privacy Protect Model Based on Positioning Middleware among the Internet of Things. In *Computer Science and Electronics Engineering (ICCSEE), 2012 International Conference on*. pp. 288–291.

LO, A., LU, W., JACOBSSON, M., PRASAD, V. AND NIEMEGEERS, I., 2006. 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*. pp. 1–8.

MACDONELL, S., SHEPPERD, M., KITCHENHAM, B. AND MENDES, E., 2010. How Reliable Are Systematic Reviews in Empirical Software Engineering? *Software Engineering, IEEE Transactions on*, 36(5), pp.676–687.

MACHARA, S., CHABRIDON, S. AND TACONET, C., 2013. 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)*. pp. 557–562.

MAYER, S., BECKEL, C., SCHEIDEGGER, B., BARTHELS, C. AND SÖRÖS, G., 2012. Demo: Uncovering Device Whispers in Smart Homes. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia*. Ulm, Germany: ACM, pp. 56:1–56:3.

MAYER-SCHÖNBERGER, V. AND CUKIER, K., 2013. *Big Data: A Revolution That Will Transform How We Live, Work, and Think*, Eamon Dolan/Houghton Mifflin Harcourt.

MEINGAST, M., KING, J. AND MULLIGAN, D.K., 2007. Embedded RFID and Everyday Things: A Case Study of the Security and Privacy Risks of the U.S. e-Passport. In RFID, 2007. IEEE International Conference on. pp. 7–14.

METZ, R., 2013. Home Tweet Home: A House with Its Own Voice on Twitter. *MIT Technology Review*.

OHM, P., 2010. Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization. *UCLA Law Review*, 57(6).

OLESHCHUK, V., 2009. 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*. pp. 336–340.

PERRY, T., Hockey Fans to Test Facial Recognition Technology - IEEE Spectrum.

PRICE, B.A., ADAM, K. AND NUSEIBEH, B., 2005. Keeping ubiquitous computing to yourself: A practical model for user control of privacy. *International Journal of Human-Computer Studies*, 63(1-2), pp.228–253.

PRIVAT, G., 2012. *Extending the Internet of Things*, Rochester, NY: Social Science Research Network.

QIN, E., ZHANG, Y.L.C. AND HUANG, L., 2013. LNCS 8017 - Cloud Computing and the Internet of Things: Technology Innovation in Automobile Service. , pp.1–8.

REAY, I.K., BEATTY, P., DICK, S. AND MILLER, J., 2007. A Survey and Analysis of the P3P Protocol's Agents, Adoption, Maintenance, and Future. *Dependable and Secure Computing, IEEE Transactions on*, 4(2), pp.151–164.

ROLFE, P., Facial recognition technology to help combat sport troublemakers. *HeraldSun*.

SAMSUNG, SAMSUNG Unveils New Era of Smart Home at CES 2014. *Samsung Electronics America*.

SANTUCCI, G., 2009. From Internet of Data to Internet of Things. In *International Conference on Future Trends of the Internet*.

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

SINGEL, R., 2004. American Passports to Get Chipped. *WIRED*.

SOGHOIAN, C., 2009. Caught in the Cloud: Privacy, Encryption, and Government Back Doors in the Web 2.0 Era. *J. on Telecomm. and High Tech. L.*, 359.

SPIEKERMANN, S., 2010. *About the "Idea of Man" in System Design - An Enlightened Version of the Internet of Things?*, Rochester, NY: Social Science Research Network.

STRAHILEVITZ, L., 2008. Reputation Nation: Law in an Era of Ubiquitous Personal Information. *Northwestern University Law Review*, 102.

SWEENEY, L., 2002. K-Anonymity: A Model for Protecting Privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(05), pp.557–570.

TAO, H. AND PEIRAN, W., 2010. Preference-Based Privacy Protection Mechanism for the Internet of Things. In *Information Science and Engineering (ISISE), 2010 International Symposium on*. pp. 531–534.

TENNENHOUSE, D., 2000. Proactive Computing. *Commun. ACM*, 43(5), pp.43–50.

UKIL, A., BANDYOPADHYAY, S., JOSEPH, J., BANAHATTI, V. AND LODHA, S., 2012. 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: ACM, pp. 75–84.

WANG, Y., More People Have Cell Phones Than Toilets, U.N. Study Shows. *TIME*.

WANG, Y. AND WEN, Q., 2011. A Privacy Enhanced DNS Scheme for the Internet of Things. In *Communication Technology and Application (ICCTA 2011), IET International Conference on*. pp. 699–702.

WEBER, R.H., 2009. Internet of Things--Need for a New Legal Environment? *Computer law & security review*, 25(6), pp.522–527.

WEBER, R.H., 2010. Internet of Things--New Security and Privacy Challenges. *Computer Law & Security Review*, 26(1), pp.23–30.

WOHLSEN, M., 2014. What Google Really Gets Out of Buying Nest for $3.2 Billion. *WIRED*.

YICK, J., MUKHERJEE, B. AND GHOSAL, D., 2008. Wireless Sensor Network Survey. *Computer networks*, 52(12), pp.2292–2330.

YU, T., LI, N. AND ANTÓN, A.I., 2004. A Formal Semantics for P3P. In *Proceedings of the 2004 workshop on*

Secure web service. ACM, pp. 1–8.

ZHU, W., YU, J. AND WANG, T., 2012. A Security and Privacy Model for Mobile RFID Systems in the Internet of Things. In *Communication Technology (ICCT), 2012 IEEE 14th International Conference on*. pp. 726–732.

Received month20XX; revisedmonth20xx; accepted month20xx