Security and Privacy Grand Challenges for the Internet of Things

Glenn A. Fink, Dimitri V. Zarzhitsky, Thomas E. Carroll, and Ethan D. Farquhar
National Security Directorate
Pacific Northwest National Laboratory
Richland, Washington, USA

Abstract — **The growth of the Internet of Things (IoT) is driven by market pressures, and while security is being considered, the relationship between the unintended consequences of billions of such devices connecting to the Internet cannot be described with existing mathematical methods. The possibilities for unintended surveillance through lifestyle analysis, unauthorized access to information, and new attack vectors will continue to increase by 2020, when up to 50 billion devices may be connected. This paper discusses various kinds of vulnerabilities that can be expected to arise, and presents a research agenda for mitigating the worst of the impacts. We hope to draw research attention to the potential dangers of IoT so that many of these problems can be avoided.**

Keywords — Internet of Things; security; privacy; grand challenge; Internet; social impact of computing

I. Introduction: The change IoT is bringing

The next five years (from 2015 to 2020) are predicted to usher in the biggest growth of the Internet yet. Gartner predicts 25 billion devices will be connected to the Internet by 2020 [25]. Cisco predicted 50 billion connected devices, including a quarter billion Internet-enabled vehicles [26]. The Internet of Things (IoT) makes the Internet pervasive and invisible, from the clothes we wear to our connected homes and vehicles to our social and work lives. From one perspective, the rise of the IoT is no change at all. It is based on the same technologies, uses many of the same protocols, and merely extends existing applications of Internet technologies. From another perspective, the IoT is changing society more fundamentally than the Internet itself did. This paper examines some of the security and privacy grand challenges that machine-to-machine communication and smart service applications present. We discuss each area and its challenges, and propose a research agenda to address the challenges of the new world these communicating devices will create for us. The paper begins by providing a background on key concepts, definitions, and areas of impact in Section II. Then, Section III uses specific examples of emerging IoT technologies to identify problems and issues in relation to the key areas we identified earlier. Section IV presents our view of the critical IoT challenges as motivated by our analysis of the state-of-the-art in commercial and research arenas. Section V summarizes the discussion and suggests directions for future research and development efforts.

II. Overview: Problems and Opportunities

New implications for society: The Internet was designed to allow people to communicate. The IoT takes people out of the loop letting machines talk unimpeded. But taking humans completely out of the loop with machines sensing, deciding, and acting on their own brings a host of potential control and monitoring problems. The invisible impact of machine decisions that affect people is definitely increasing [1]. Crime may attain staggering new dimensions [2]. Currency may become invisible and harder to track, and consumers will be even more barraged with advertisement messaging when thousands of devices betray their every interest and movement. Clearly, IoT is changing society fundamentally.

Redefinition of Identity: Identity, once solely a property of living creatures, used to be traced to a physical body with distinctive characteristics like fingerprints. But with the rise of IoT the dominant population of identities on the planet has become that of machines, with 10 billion devices already on the Internet and exponentially more each year. Each of these entities has authority to access some service (whether wireless networks, bank accounts, or control of physical equipment), and access is granted by identity demonstrated through possession of credentials. Unlike physical humans who are difficult to clone, identities in cyber space are very easy to spoof [3]. Physically Uncloneable Functions (PUFs) show promise to enable real identity for hardware devices. PUFs are essentially a biometrics analogue for hardware. Real progress has been made in recent years [24], but electronic boards can be swapped, firmware can be overwritten, and software can be changed [18]. Whether PUFs alone can secure hardware is uncertain. How can we presume the durability of an identity once it is trusted? Given the fluid nature of identity in the IoT, understanding and management of this trust becomes important [19].

Devices may be replaced frequently (*e.g.*, cell phones), relationships may be ephemeral (*e.g.*, a purchase from a vending machine), and trust can have transitivity (as when devices owned by friends are also to be trusted). The sheer number of devices expected to join the IoT will make any but the most efficient trust-management algorithms useless. Widely used public-key infrastructure (PKI) and the X.500 framework of certification authorities and revocation lists are completely unworkable with such huge and dynamic populations. Alternative, lightweight trust models such as Simple Public Key Infrastructure (SPKI) may deserve a second look in the IoT context. Webs of trust will become more complex, and many more levels of trust within communities may be required.

Although "security through obscurity" is generally rejected as a security strategy, obscurity may be a tenable approach to

978-1-4673-7648-8/15/$31.00 ©2015 IEEE

security for power and storage deprived devices that cannot be defended any other way. With many devices (like sensor networks in connected automobiles) clear lines of ownership, responsibility, and control are difficult to define. Is the manufacturer, operator, or owner responsible for the safety of these systems? Who legally owns the data they generate? Finally, the implications of deception and the difficulty of attribution in the IoT world present daunting challenges of their own.

Architectural ambiguities: Internet Protocol version 6 (IPv6) laid the groundwork to have more device addresses than there are baryonic atoms in the known universe. But can the Internet itself support the actual existence of extremely large numbers of devices in complex and dynamic relationships with one another? IPv6 (via 6LoWPAN) is just one of many possible transport layers in use with IoT devices—there are dozens if not hundreds of different protocols at different communication stack layers [20]. Many cheap and lightweight IoT devices may implement micro-protocol networking stacks that may not be faithful implementations of the standards. The problem is exacerbated when devices of different types try to communicate. Thus, the IoT is actually millions of micronets with varying kinds of Internet gateways rather than an extension of the existing Internet. The IoT currently has more in common with industrial control systems (ICS) than with a true Internet.

Gartner predicts that no dominant IoT platform will exist through at least 2018; applications will remain a mash-up of semi-congruent parts from various providers [25]. Until a unified ecosystem emerges, there will be no coherent set of business or technical models for the IoT. Because standards are new and untried, most IoT projects will need to invent custom elements from various technology service providers in the IoT with no clear choice. History suggests that when a disruptive technology appears, the majority of *de facto* ecosystems will fail during the working lifetime of current projects, requiring system integrators to devise careful strategies to future-proof their work. It is reasonable to assume that pervasive flaws will persist in a system this complicated for decades. In the next sections, we will examine these challenge areas in more detail.

III. Societal Effects of IoT

"The Internet is the first thing that humanity has built that humanity doesn't understand, the largest experiment in anarchy that we have ever had." — Eric Schmidt [27]

A. Standards

Standards exist or are being defined for the entire network stack of IoT [6] from the physical and medium access layers (IEEE 802.15.4, IEEE 802.15.4e, and 6LoWPAN), to the routing layer (ROLL RPL and IPv6), to the application layer (Constrained Application Protocol, CoAP) [6]. This suite of protocols, however, represents only the most standardized of all choices that can be made. Thousands of interoperable protocols can be derived from the standards, each differing according to manufacturer choices. However, subtle differences between protocols are excellent places to find cyber exploits.

Unfortunately, nascent IoT designs will immediately face a highly polished threat from the adoption of well-practiced attack techniques. Ptacek and Newsham [4] showed how insertion, evasion, and denial-of-service approaches could be used to carefully target individual parts of a security system and disable them. Insertion and evasion attacks exploit the differences between intrusion detection systems (IDS) and the operational systems they protect. By inserting packets that the target will accept but the IDS rejects as bogus into an attack stream, the attacker evades the IDS and succeeds in hitting the target with an exploit the IDS never sees. Similarly, by injecting packets that the target will reject but the IDS will accept, the attacker injects distractions into the IDS's view of the situation while the end system is affected.

In the IoT, devices and applications are often chained end-to-end in a mash-up of existing services. If data can be inserted that is illegal to the protocols of all but one kind of devices in the chain, an attacker can target that kind of device. Conversely, an attacker could use such differences to prevent just one type of device from seeing an attack. For example, if the packet fragmentation and reassembly logic of two nearly compatible implementations of 6LoWPAN differs, an attacker may be able to send something syntactically legal but semantically disruptive to the receiver. Attackers thrive on taking advantage of the resulting undefined states to disable security monitoring, escalate privileges, or simply deny service. Exactly what the outcome will be depends on the desires of the attacker and the ambiguities between protocol implementations.

Thus, while standard protocols ensure system security by definition, they cannot promise to protect unexpected combinations of particular implementations of these standards. Although there are encryption, integrity, and authentication methods described for each protocol level in the stack, most standards have minimal adherence guidelines that include no security as an option. If history is any predictor of future tendencies, key management will be the dominant expense in any encryption implementation in the IoT. Much encryption will probably be done with the simplest sort of key exchange, pre-shared secrets. Once these static keys are discovered, it will be simple to implement spoofing attacks. The standards-implementation morass will host many types of attacks that will lead to many more security and privacy challenges for society.

The success of the IoT depends on the existence and well-defined interoperability of global standards. IoT is designed to simplify information exchange, but the underlying infrastructure is very complex. The *prima-facia* simplicity of a seamless IoT experience hides a highly nuanced architecture. Individual decisions on storage, messaging and routing protocols, security, directories, analysis, automation, APIs, and many other elements make agreement on standards difficult [28]. Such technological complexities may prevent IoT from reaching the optimistic predictions currently being made.

B. Privacy

Many devices already on the Internet are related to critical infrastructures. As more parts of our lives are connected, more of the Internet becomes that critical infrastructure. "Big Data, together with the Internet of Everything, will expand existing and create new types of critical infrastructure. This in turn will create new privacy issues as these categories of data and their option value will offer new insights." [2]

Voluntary self-publication has already reached amazing proportions with blogs and social networks. But the involuntary gathering of data from a nearly ubiquitous collection of location, audio, video, and other sensors is likely to increase as Internet-connected sensing devices become commonplace. Without technological and regulatory protections, ubiquitous interactions with devices, components, resources, and services will generate voluminous data that can be used to identify and track individuals and their behaviors. Recent news demonstrates the threat of (even anonymized) "Big Data" to individuals' privacy. [8] In this section, we discuss the privacy impacts to society of these data streams enabled by IoT and the grand challenges that arise from them.

1) Location data

In a market-based, free-enterprise society, getting customers close enough to the product to be enticed into buying it is critical to sales. Additionally, knowing the location patterns of target demographics helps advertisers plan effective campaigns due to better determination of the desires of their customers.

IoT is already enabling social networks to take on a physical dimension, via proximity. However, since about 2009, proximity-based social networks have dwindled into niche applications such as Google's Waze (www.waze.com), a community-based traffic navigation application. People seem hesitant to share their location data intentionally, but IoT sensors owned by others may make anonymous sharing implicit. Consider that only four location data points are required to re-identify 95% of persons. [7] Location data has thus become a *de facto* identifier. It may be considered a commodity for commerce and governments, and it is likely to become an enabler for targeted crime.

2) Audio data

Recently, the BBC reported Samsung's disclosure that their Smart TV's voice activation feature "listens" to what people in its proximity say, and it may share that information with the manufacturer or with third parties [29]. Apple's Siri has long had the "Hey Siri" feature that listens to audio when the feature is turned on and the phone is plugged in. Not surprisingly, many devices around the home (dishwashers, washing machines, coffee makers, entertainment systems, etc.) may also gain this voice activation feature. Voice data must be continuously collected and uploaded because the device cannot tell when something spoken will be a command. The convenience of voice commands thus implies continuous, ubiquitous audio collection. Additionally, it may be advantageous for devices to isolate voices from one another and potentially understand which individual is speaking. Commands may be expressed in a variety of ways potentially requiring a large corpus of spoken data to account for dialect and pronunciation differences. This corpus may be collected and maintained by the device manufacturer, the device itself, or even a third-party provider such as Nuance, the vendor used by Samsung and many other companies, including auto manufacturers.

Unfortunately, it appears to be accepted practice to require users to agree to terms of service that force audio collection and analysis features on them. But less acceptable practices also happen: LG was found to not only send information when the collection feature was off, but their televisions were also collecting and sending back information from privately owned Universal Serial Bus (USB) drives plugged into the sets [30]. Since it is not immediately apparent when audio is collected and uploaded, breaches of privacy are difficult to avoid.

3) Video data

Personal drones already accompany amateur athletes as selfie-drones and personal trainers. But the video they generate and upload can be used for a variety of other purposes. Continuous video recording and posting to the Internet is becoming an accepted part of life. However, the legal uses of such video have only begun to be examined in the courts [31]. Currently, more than 300 hours of video are uploaded to YouTube each minute [32]. Children already grow up with Internet-enabled devices in hand. Each generation is more recorded than all previous ones. Within a few years, personalized movies like Boyhood could augment yearbooks and photo albums. For example, in early 2014 Facebook automatically generated "hundreds of millions" of personalized look back videos for its users to mark the company's tenth-year anniversary [33].

Ubiquitous video collection shares many of the concerns that audio collection but also adds others. Audio is mostly collected to process commands, not to gather message information, but video is often gathered to tell stories. Frequently, it is uploaded and stored online permanently. Video also carries many highly identifiable individual features and metadata that can make identification and placement of persons possible. Automated means of face recognition has been around for well over a decade now [16] and is only improving. Thus, captured video becomes another part of a person's identity.

4) Digital Identity

The current best practice for providing identity for digital devices is via pre-shared secret keys stored in nonvolatile read-only memory and used to create digital signatures or to encrypt data. This approach is expensive, vulnerable to tampering or side-channel attacks, and requires continually powered active tamper detection/prevention circuitry. Alternatively, PUFs [18], [24] may form the basis for device authentication (see Figure 1 below). Small, unavoidable manufacturing differences give each integrated circuit slightly different performance characteristics. These can be combined via on-chip logic to produce unique, extremely difficult to forge, response bit streams when the device is challenged for its identity. However, when devices have multiple boards or component parts, PUFs would have to be combined to ensure the *whole* device was the same.

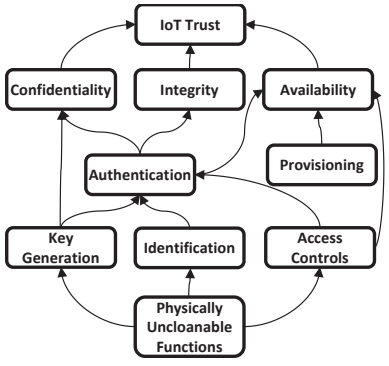

Figure 1. Prerequisite dependencies for establishing IoT trust.

Security of users depends on much more than assurance of device identity. Much can be hidden in the multiple layers of software interposed between the user and the hardware. Even perfectly unforgeable device identifiers cannot provide full assurance that a man-in-the-middle attack is not used to pass along identifiers and subsequently fool authentication systems.

5) Vehicles

Most new cars come with wireless features and ability to gather, store, and transmit drivers' data (including location), but a report from the US Senator Edward Markey's office says the protection afforded these systems is "inconsistent and haphazard" with data often being transmitted insecurely [34]. Many vehicles collect and transmit data on vehicle performance, driving history, navigation, and last parking location. Manufacturers have actually used this data to rebut negative press reports on test drive experiences. Remote disabling of cars and navigation systems can be used by manufacturers in case an automobile is stolen or owners default on their loan payments. When such capabilities exist, ownership of the vehicle is in question.

These capabilities would simply provide an enhanced user experience if they were not also avenues for attack. Electronic Control Units (ECUs) are embedded computers that control automotive subsystems for major functions of modern cars. In newer models, it is common to have thirty or more ECUs all connected to intra-car networks. Because they were originally part of a closed network that required physical access (often through partial disassembly of the automobile) these systems were not hardened to implement even the basic aspects of security like confidentiality, authentication, and integrity [13]. Most new models of cars also communicate with the outside world through a variety of external protocols such as the USB, Bluetooth, wireless (Wi-Fi) or cellular (3/4G) networks. The mappings between these external points of connection and the vulnerable internal critical systems is often not well known, and researchers have demonstrated numerous attack opportunities and actual exploits that can be conducted ranging from annoyances to life-threatening severity. Studnia, et al. [13] present a taxonomy of attacks against connected vehicles including motivations such as theft, unauthorized modification, sabotage, intellectual property theft, and privacy breaches.

We expect a trend of increasing external connectivity as both car-to-infrastructure and car-to-car networking applications arise. We also expect it will be difficult to harden the internal networks that depend on static trust relationships and to understand fully how to separate them from external networks. This is especially true since some applications such as creating automatic "platoons" of self-driven cars [14] requires external access to critical systems like steering, brakes, and acceleration. Finally, the lifecycle of automobiles is around twenty years, much longer than for workstations with similar computing power. Thus, automotive hardware must remain securable via software upgrades for much longer than most commercial operating systems are maintained. This long-term maintenance is likely to be a significant expense for manufacturers.

6) Other personal data

Personal health-monitoring devices already track health data and increasingly publish it to the Internet for consumption by doctors, Facebook friends, etc. Anthropometric data (gait and gesture identification, etc.) can be gathered from data such as Kinect video from toys or gesture-based user interfaces [15]. These developments increase convenience and may help IoT-type devices recognize the circumstances of their use (running, in a quiet theater, etc.) so that they may adjust their behavior accordingly. However, if this data is shared it too will become part of a person's identity. [21]

Life-pattern analysis can benefit both consumers and advertisers by attracting buyers to the products and services they desire. As IoT-enabled devices become ubiquitous, their embedded sensors will enable new streams of data for life-pattern analysis and new venues for advertisement. Just as independent bloggers receive compensation for allowing advertisements on their blogs, people with wearable electronics may soon earn money by becoming walking digital billboards via invisible, near-field transmissions that advertise their sponsors' products.

7) Ubiquitous sensing:

Governments and corporations are seeing the IoT's tremendous capacity as a global sensor network. Mostly, this is an attempt to fulfill the current roles of government better through better sensing. Detecting illicit trafficking in drugs, nuclear materials, or abducted persons may be easier if the commodities are electronically tagged to make identification and tracking easier. Combining sensor input from many vantage points can help pinpoint terrorist activity before it happens.

Singapore is perhaps the best-known example of a state that uses an IoT infrastructure extensively to protect its populace [35]. Most Singaporeans perceive the widespread surveillance their government has implemented as a safeguard. However, political dissidents, including cartoonist Chew Peng Ee, have objected to the government's broad use of power and overly general interpretation of the law. Most citizens see the government as protecting them against the potential data abuses they fear private companies might commit. Americans, on the other hand, tend to trust corporations more than their government. The Snowden leaks produced a public outcry against invasive government but no significant policy change ensued. It may be that citizens are coming to the sad conclusion that, as Vint Cerf said, "Privacy may actually be an anomaly." In the face of widespread surveillance by state and non-state actors, freedom of expression is endangered, even if unintentionally. If all the conveniences IoT devices provide ultimately result in loss of basic freedoms enjoyed by more primitive societies, will we still consider the trade advantageous?

C. Security

Security on the Internet has frequently been relegated below the primary task of creating marketable features. However, given the increasing levels of crime associated with the growth of the Internet and the projected growth of the IoT, security should be of great concern.

1) Crime

The IoT will likely expand criminal uses of the Internet simply by providing vastly more devices for criminals to exploit and multitudes of new protocols to obfuscate their trail. The Internet has already been exploited by crime organizations to amplify the impact of their activities. Many criminals have

found ways to make money by selling botnets, attack kits, and hacking services. Anonymous crime-as-a-service allows one to commit crimes from a great distance and be protected by poor traceability and the incompatibilities of international laws.

Another problem is the very invisibility of ubiquitous computing itself, "One of the threats arising from [the Internet of Things] is that, whereas people often consciously log in to computers and even smartphones, they may not be aware of how they are connected to the IoT environment." [2] Wearable computing and personal area networks make it possible to infer much about a person's location, surroundings, preferences, and life patterns that help criminals take advantage of victims.

Near-field communication (NFC) is making physical possession of credit cards less important, just as credit cards have eliminated much of the need for cash. Money is increasingly becoming a virtual construct that can be whisked off with the transmission of a few bits. While NFC is protected by biometrics, biometric identification is also susceptible to forgery, as recently demonstrated by Jan Krissler, who was able to steal the thumbprint of the German Minister of Defense using only publically available Internet video [36].

A multiplicity of devices provides more hosts with fewer people monitoring them. Embedded devices are often difficult to keep up to date with firmware, and commodity operating systems change so quickly that the expected upgradeable lifetime of such devices is less than ten years. After a while, they simply cannot receive updates because of hardware incompatibilities, and these devices may become vulnerable to attack. This is already happening in expensive legacy equipment such as mass spectrometers and medical equipment. Embedding computation in massive arrays of vulnerable Internet-connected devices could allow for the creation of botnets of billions of devices all over the world.

Additionally devices may have multiple identities. Reputation-based trust mechanisms are unenforceable when anonymous and pseudonymous identities cannot be linked to the real identity of a device or component. The ability of an entity to assume an unlimited number of identities is referred to as the *"Sybil attack"* in the computer security literature [3]. The question becomes, who will mediate the creation and management of these identities? If attestations to actual identities are made through a trusted third party, trust may be established only to the degree that the third party is indeed trustworthy.

2) Cyber warfare

IoT is often associated with cyber-physical systems, machines with the ability to manipulate their physical environment that are controlled by embedded computer systems and connected to the Internet. These devices include things like supervisory control and data acquisition (SCADA) systems and industrial control systems (ICS). Securing these devices and understanding their connectivity and remotely accessible capabilities is crucial to future societal security. For example, a 2014 Bloomberg report found that a 2008 explosion in a highly secure Turkish pipeline was caused by computer sabotage:

"Hackers had shut down alarms, cut off communications and super-pressurized the crude oil in the line. The main weapon at valve station 30 on Aug. 5, 2008, was a keyboard." [12]

This pipeline had been considered the most secure in the world with concrete housings protecting miles of pipes and surveillance cameras covering its entire length. However, the attackers subverted the security camera network to become their point of entry into the system. From there they scrambled the instructions that regulated pressure in the pipelines creating a huge and costly explosion. If nation-state actors were involved, this attack would be the first instance of kinetic cyber warfare, two years before the widely reported StuxNet infestation was discovered in Iran.

3) Emergent Behaviors

Emergent behaviors are potentially unexpected side effects that happen in complex systems usually because of feedback relationships. Ideally, emergent behaviors should be allowed for in the system design as has been done in the electric power grid, but one can only design around feedback relationships that are understood. With electric power, blackouts and cascading failures are fairly well understood, so resilience has been built into the system to account for these emergent behaviors. Internet Transmission-Control Protocol/Internet Protocol (TCP/IP) is another example of a carefully designed protocol suite that prevents traffic flow anomalies like congestion and dropped connections from becoming negative emergent behaviors. However TCP/IP's approach (throttling back data rates) works poorly for Wi-Fi networks where packet loss may be due to weak signal, not traffic congestion. TCP's approach can create a negative emergent behavior in wireless networks that further harms already bad throughput. Many devices that use these same protocols for wired and wireless connections already populate the IoT, and these devices are widely dispersed, massively heterogeneous, semi-autonomous, and often unreliable. The possibility of unintended negative emergent behaviors may be very large and significant.

IV. A Grand Challenges Research Agenda for IoT

To meet some of these many challenges we recommend research in a variety of areas that will contribute to solutions.

A. Scientific and Technical Challenges

Back-of-the-envelope calculations of data transmission rates predict significant engineering challenges. Current estimates on the amount of data transmitted on the global Internet (wired and unwired) vary, but six exabytes (6×10^{18} bytes) per month seems to be a reasonable estimate. Assuming the number of devices grows at the projected rates, and given a packet size of 100 bytes at 1 Hz collection rates, the IoT will require the capacity to transmit six exabytes every *second* by 2020. This rough figure ignores collisions, errors, and other such events that necessitate re-transmission.

The obvious need is for fundamental research in edge analytics to reduce the amount of data that must cross backbone network segments. Another clear need is the ability to store this massive amount of data. Even if we are able to reduce the storage requirements dramatically, storage of even a tiny fraction of the available data will require high-bandwidth, continuous access to high-volume storage.

Since data collection is often invisible to those it affects, an area of potential research may involve enhancing user awareness of audio and video collection. This may involve standards

that include user-notification and data tagging to make recording more detectable. We also need a general estimation science to approximate how difficult it is to identify an individual from a set of data sources that contain clues about his or her identity.

Technical challenges in the automotive domain will require reimplementation of many known kinds of fixes to new hardware in a new domain. Adding encryption and protocol hardening to low power, real-time automotive control networks will require extensive (and expensive) re-engineering of existing systems. Certifying the security of these open systems will be no easier than it was for workstations industrial control systems. Beyond encryption, [13] suggests that anomaly detection and embedded software integrity checks are major research challenges. Subsets of many of these problems have been solved before elsewhere, but every solution requires engineering tradeoffs and may produce unintended side effects.

Because IoT applications will compose services from multiple machines we must be able to understand how security and privacy functions may be composed across systems and protocol layers. For example, if an Internet-connected car synchronizes its owner's private contacts from his personal phone, how do we ensure that the data retains the same privacy protections in the car as on the phone? We must be able to extend standard security models such as Mandatory Access Control (MAC) to systems of systems where data is portable but must still be securely accessed. Data must be secure from creation to destruction, whenever it is transmitted or stored. As far as we know, Ionic Security is the only company that claims to be able to do this, and its product has not yet been tested. In general, if such data security is possible on limited IoT devices, then a key step will have been taken for privacy and security on the Internet.

One promising area of future research is in biologically inspired approaches [17]. Animal brains are well suited for dealing with real-world ambiguities and highly dynamic environments. Multiplicity, cooperation, and wide geographic distribution are crucial in ensuring stable biological communities, and we expect that technological approaches inspired by biological systems will be effective in addressing some of the IoT challenges we have identified. We are already starting to see commercial solutions offering self-healing properties for networking applications, such as cloud based elasticity, virtualization, and multiple-zone data center availability with instant failover.

Simplified trust models such as SPKI are good candidates for IoT. The Common Criteria [9] has defined a set of privacy qualities for privacy-preserving identity management:

- **Anonymity:** An individual may use a resource without disclosing identity.
- **Pseudonymity:** An individual may use a resource without disclosing identity, but remain accountable for use.
- **Unlinkability:** An individual may make multiple uses of a resource without others being able to link these uses together.
- **Unobservability:** An individual may use a resource without others, especially third parties, being able to observe that the resources are being used.

Reflection on the above challenges suggests that we may require anonymous and pseudoanonymous identities to evolve with spatiotemporal variables. Neither nearby vending machines nor passing vehicles should be able to digitally determine that repeated interactions are the same individual or enable other entities to ascertain that identity. Identities should be inexpensive to obtain and replace, and a single entity or authority should not control the allocation of these privileges.

We also envision a subset of purpose-built, certified, federated technologies intended to be deployed to prevent individuals from accidentally (or purposefully) opting out of the safeguards. Biometrics alone are not the answer. Instead, secure systems should rely on cryptographic means to incorporate random noise to "whiten" the biometric artifact, thwarting unauthorized duplication. To do this we require a new mathematical approach to describe security and privacy implications of "connecting anything to anyone at any time." We need a closed form expression of our degree of certainty that information is secure, privacy is upheld, and systems will function properly in an IoT environment. Standards enable interoperability, not security. Security forms a tight bound on interoperability limiting possibilities to what is desirable by device and data owners.

We must emphasize that securing the IoT, from a technical standpoint, requires end-to-end solutions, not a collection of point solutions. In the previous example of PUFs we showed that even unforgeable device identifiers are only part of a larger security system that involves other layers of software and hardware. Thus, another grand challenge in IoT is creating a standard security stack similar to the network stack, with standard interfaces and degrees of assurance. The micronets of IoT devices must be integrated into a higher-level Web of Things (WoT) using semantic services and standard resource description ontologies [23]. This will give machine-to-machine (M2M) devices a standard interface that people will be able to interact with trust. Conceptually, the WoT is where the IoT meets the traditional Internet and the cloud [5].

B. Social and Regulatory Challenges

Massive data from ubiquitous sensing applications presents the grand social challenge of IoT: what may be done with the data? The data is being gathered continuously, it might never be erased, and individuals have no control over its protection, dissemination, use, or inferences drawn from it. Because personal data may be collected on persons from various nationalities, several states are already requiring that personally identifiable data concerning their citizens (e.g., shopping carts, banking information, address books, etc.) must be stored on devices that are physically within their countries' borders. Storage must have a physical location, and it makes sense that data should be stored topologically near devices if possible. But the recent "right to be forgotten" enacted by the European Union governments takes on new dimensions when IoT is considered. This is especially true when we consider the entire device and its data lifecycle, including transfer of ownership and eventual disposal. Broad device categories based on physical size, the type of power source, typical use patterns, and so on should be established so that manufacturers and consumers can agree on the expected levels of support. Comprehensive regulatory protections and technologies to support them must be created

and enabled early on in the device's (and its software) manufacture in order to properly address these grand challenges.

Looking at age-related demographics for users of social networking sites like YouTube, Twitter, and Facebook, it appears that today's youth are less concerned with privacy than their elders. It remains to be seen whether a person's privacy preference changes with age, but it is important for the technology to be responsive to such gradual shifts. New regulations and supporting technologies that can safeguard individuals' privacy from misuse over long periods of time are needed. Additionally, ways for individuals to query massive data for information that pertains to them must be enabled. Finally, individuals need the ability to change incorrect conclusions inferred about them by ubiquitous sensing systems.

End-User License Agreements (EULAs) often force users into accepting compromise of their privacy in order to use features (like voice recognition) of devices they have purchased. Removal of forced EULAs that protect such features will restore freedom to the people whose homes, vehicles, clothes, etc. are continuously recording and transmitting personal data.

Today, we may be unable to predict all the ways an adversary could exploit the kinds of information that IoT could provide in the future. Criminals will exploit it, but we would like to be able to make informed engineering trade-offs that decide what will be more susceptible to exploitation so we can prepare appropriately for the consequences. Much may be done here both technically (e.g. data encryption in flight and at rest), and legally, considering what someone might do with this information. Additionally, the psychological implications of these systems should be considered. What do breaches in security do to the victims involved? What can be done to encourage good security practices? How can we change people's behaviors to enhance security further?

One means of inducing security-conscious behavior is through incentives, such as a variation on the "pay to play escrow" approach that motivates users to act securely based on their own investment in the system [37]. However, the incentives must be carefully structured to avoid adverse effects by borrowing from economics, game theory, and distributed mechanism design (see [10], [11] for suggested methods).

A significant portion of IoT devices currently on the market feature restricted user interfaces with various physical constraints and reduced customization. Lack of consistency, inflexible user interfaces, and complex interactions between systems create unique challenges for implementing assistive technologies for users with disabilities [22]. Invisible machine-to-machine (M2M) interactions hinder users' perception of the risks inherent in connecting these systems to the Internet or other networks. Going back to our earlier example of the cell phone pairing with the automobile, when the phone shares its contact list with the vehicle's on-board computer, the safeguards used to protect the information are invisible to the user. Once shared, the user might not be able to tell when the car stores this information locally (hopefully in an encrypted format) and whether it shares the data with nearby vehicles on the road. Standardized laws of data management must be designed and written into standard protocol stacks in a manner reminiscent of Asimov's three laws of robotics. For instance, a device should not share information stored in it with anyone but the owner of the data, unless specifically instructed to share such data. Unfortunately, as envisioned, Asimov's laws entail a certain awareness that IoT devices are unlikely to have, thus any such future "laws" must be mathematically defined, so that they work for the majority of use cases. Asimov's own stories are replete with examples of unexpected behaviors that arose from the real-world interactions of seemingly intuitive laws. We expect that defining guidelines for well-behaved IoT implementations will be comparably complex and challenging.

V. Conclusion

In this paper, we addressed a range of technical and social concerns specific to the rapid growth of IoT technologies being introduced into the market. Existing vulnerabilities of Internet protocols and lack of sufficiently powerful mathematical analysis tools cause us to anticipate a rapidly growing set of challenges associated with the early adoption of IoT systems. Social impact of these technologies is far-reaching, and will require involvement and coordination between government regulatory agencies, private industry, academia, and international standards bodies. We have cited the necessity of sustained research into novel approaches to encryption and authentication to address the massive scale of IoT networks, together with theoretical methods for incorporating low-power requirements of embedded devices. Increased use of scalable, non-federated authentication technologies is an important enabling factor, especially when supplemented by innovative uses of hard-to-duplicate physical attributes. Emerging research trends indicate that biologically inspired approaches that exploit our understanding of natural self-organization and energy optimization provide unique advantages in dealing with such resource-constrained domains.

IoT ecosystems should provide incentives for users to be proactive in securing their personal data. Open questions regarding device data ownership and the gradual change in perceived security benefits and privacy concerns over the owner's lifetime necessitate a comprehensive lifecycle model covering the deployment, maintenance, and eventual retirement of IoT systems. Government support in establishing global security standards is important to establish trust in the growing network and the data it contains. An international consensus is needed to prevent fragmentation of various privacy initiatives, and to ensure consistent level of protection irrespective of the device or its data storage geographic location. Of special interest are M2M interfaces that obscure transfers of personally identifiable data, indicating a need to provide the means by which individuals can assess and control data that has been compiled on their behavior. We also recognize the tremendous potential of IoT medical devices, and in particular, the benefits for those affected by disabilities. At the same time, the potential for misuse of such sensitive information means that we need to focus on new approaches to identity and trust management that are able to scale with the growing IoT.

Acknowledgement

The authors would like to acknowledge Pacific Northwest National Laboratory's leadership enabling this research, which corresponds to PNNL report number: PNNL-SA-108586.

REFERENCES

[1] A. Kirilenko, A. S. Kyle, M. Samadi, T. Tuzun, "The flash crash: the impact of high frequency trading on an electronic market," September 2014. Available at http://dx.doi.org/10.2139/ssrn.1686004.

[2] European Cyber Crime Center, European Police Office, "EuroPol Internet organized crime threat assessment," 2014. Available at http://www.europol.europa.eu.

[3] J. R. Douceur, "The Sybil attack," First International Workshop on Peer-To-Peer Systems (IPTPS '01), pp. 251-260, 2002.

[4] T. H. Ptacek and T. N. Newsham, "Insertion, evasion and denial of service: eluding network intrusion detection," Secure Networks Inc., January 1998.

[5] ITU-T SG13, "Future networks including cloud computing, mobile and next-generation networks," April 2015, http://www.itu.int/en/ITU-T/studygroups/2013-2016/13/Pages/default.aspx.

[6] J. Granjal, E. Monteiro, J. Silva, "Security for the Internet of Things: a survey of existing protocols and open research issues," IEEE Communications Surveys & Tutorials, pp. 1-1, 2015.

[7] Y-A de Montjoye, C. Hidalgo, M. Verleysen, V. Blondel, "Unique in the crowd: the privacy bounds of human mobility," Sci. Rep., vol. 3, 2013.

[8] Y-A de Montjoye, L. Radaelli, V. K. Singh, and A. Pentland, ``Unique in the shopping mall: on the reidenifiiability of credit card metadata," Science, vol 347(6221), January 2015.

[9] Common Criteria, "Common Criteria for information technology security evaluation," 2014. Available at https://www.commoncriteriaportal.org/files/ccfiles/CCPART2V3.1R4.pdf.

[10] J. Shneidman and D. C. Parkes, "Specification faithfulness in networks with rational nodes," In Proceedings of the ACM symposium on Principles of Distributed Computing (PODC '04), pp. 88-97, 2004.

[11] J. Shneidman, D. C. Parkes, and E. Massouli'e, "Faithfulness in Internet algorithms," In Proc. of the Workshop on Practice and theory of Incentives in Networked systems (PINS '04), pp. 220-227, 2004.

[12] T. Zhang, H. Antunes, S. Aggarwal, "Defending connected vehicles against malware: challenges and a solution framework," IEEE Internet of Things Journal, vol 1, no. 1, pp. 10-21, February 2014.

[13] I. Studnia, V. Nicomette, E. Alata, Y. Deswarte, M. Kaaniche, and Y. Laarouchi, "Survey on security threats and protection mechanisms in embedded automotive networks," In Proc. Of the Dependable Systems and Networks Workshop (DSN-W), pp. 1-12, June 2013.

[14] E. Coelingh and S. Solyom, "All aboard the robotic road train," IEEE Spectrum, vol. 49, no. 11, pp. 34-39, November 2012.

[15] B. C. Munsell, A. Temlyakov, C. Qu, and S. Wang, "Person identification using full-body motion and anthropometric biometrics from Kinect videos," In Proceedings of the Twelth International Conference on Computer Vision, vol. III, pp. 91-100, 2012.

[16] T. Choudhury, B. Clarkson, T. Jebara, and A. Pentland, "Multimodal person recognition using unconstrained audio and video," In Proceedings of the International Conference on Audio- and Video-Based Person Authentication, pp. 176-181, 1999.

[17] C. Liu, Y. Zhang, and H. Zhang, "A novel approach to IoT security based on immunology," In Proceedings of the International Conference on Computational Intelligence and Security (CIS), pp.771-775, 2013.

[18] T. Xu, J. B. Wendt, and M. Potkonjak, "Security of IoT systems: design challenges and opportunities," In Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD '14), pp. 417-423, 2014.

[19] M. Abomhara and G. M. Koien, "Security and privacy in the Internet of Things: current status and open Issues," In Proceedings of the International Conference on Privacy and Security in Mobile Systems (PRISMS '14), pp. 1-8, 2014.

[20] C. Hennebert and J. Dos Santos, "Security protocols and privacy issues into 6LoWPAN stack: a synthesis," IEEE Internet of Things Journal, vol. 1, no. 5, pp. 384-398, 2014.

[21] A. Grau, "Can you trust your fridge?," IEEE Spectrum, vol. 52, no. 3, pp. 50-56, 2015.

[22] M. C. Domingo, "An overview of the Internet of Things for people with disabilities," Journal of Network and Computer Applications, vol. 35, no. 2, March 2012.

[23] E. G. Stephan et al., "Semantic catalog of things, services, and data to support a wind data management facility," Information Systems Frontiers Journal, pp. 1-13, 2015.

[24] C. Herder, M. Yu, F. Koushanfar, and S. Devadas, "Physical unclonable functions and applications: a tutorial," Proceedings of the IEEE, vol. 102, no. 8, pp. 1126 – 1141, August 2014.

[25] B. Rossi, "Gartner's Internet of Things predictions," Information Age, Vitesse Media, January 2015. Available at http://www.information-age.com/technology/mobile-and-networking/123458905/gartners-internet-things-predictions.

[26] D. Evans, "The Internet of Things: how the next evolution of the Internet is changing everything," Cisco Internet Business Solutions Group, 2011. https://www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf.

[27] J. Taylor, "Google chief: my fears for generation Facebook," The Independent, Independent Print Ltd, August 2010. Available at http://www.independent.co.uk/life-style/gadgets-and-tech/news/google-chief-my-fears-for-generation-facebook-2055390.html.

[28] C. Green, "Making the Internet of Things a business reality," Information Age, Vitesse Media, November 2014. Available at http://www.information-age.com/technology/mobile-and-networking/123458674/making-internet-things-business-reality.

[29] BBC News, "Not in front of the telly: warning over 'listening' TV," February 2015. Available at http://www.bbc.com/news/technology-31296188.

[30] T. Cushing, "LG Smart TV caught collecting data on files stored on connected USB drives," Techdirt, Floor64, November 2013. Available at https://www.techdirt.com/articles/20131119/06503625288/lg-smart-tv-caught-collecting-data-files-stored-connected-usb-drives.shtml.

[31] P. G. Geraghty, "United States of America National Transportation Safety Board Office of Administrative Law Judges Decisional Order," Docket CP-217/2012EA210009, pp. 1-9, March 2014. Available at http://www.ntsb.gov/legal/alj/Documents/Pirker-CP-217.pdf.

[32] YouTube, "Statistics," Google Inc., April 2015. Available at http://www.youtube.com/yt/press/statistics.html.

[33] M. Snider, "Facebook lets users look back at decade," USA TODAY, Gannett Co., Inc., February 2014. Available at http://www.usatoday.com/story/tech/2014/02/05/facebook-movies-10th-anniversary/5235943/.

[34] E. Markey, "Tracking & hacking: security and privacy gaps put American drivers at risk," Congressional Report, Feburary 2015. Available at http://www.markey.senate.gov/imo/media/doc/2015-02-06_MarkeyReport-Tracking_Hacking_CarSecurity 2.pdf.

[35] T. Lee, "Singapore an advanced surveillance state, but citizens don't mind," Tech in Asia, November 2013. Available at https://www.techinasia.com/singapore-advanced-surveillance-state-citizens-mind/.

[36] 46halbe, "Fingerprint biometrics hacked again," Chaos Computer Club, December 2014. Available at http://www.ccc.de/en/updates/2014/ursel.

[37] Y. Ilyin, "Pay to play again: a cryptolocker variant goes after the gamers," Kaspersky Lab Business, Kaspersky Lab ZAO, March 2015. Available at http://business.kaspersky.com/pay-to-play-again-a-cryptolocker-variant-goes-after-the-gamers/3715.