A Framework for Analysis of Quotidian Exposure in an Instrumented World

Lisa A. Shay, Gregory Conti, Dominic Larkin, John Nelson
Cyber Research Center
United Stated Military Academy
West Point, NY 10996
{Lisa.Shay, Gregory.Conti, Dominic.Larkin, John.Nelson}@usma.edu

***Abstract*— For a variety of often reasonable motives such as increased security or increased profit, individuals and organizations fill our world with sensors and data collection systems that sample and track our day-to-day activities. Some people freely invite tracking into their lives. Others are enticed by offers of discounts or even free products and services. But frequently our lives are quietly sampled, unbeknownst to us, by those with the power to do so. As a result, individuals face a rapidly declining freedom to lead a private life. While significant sampling and tracking occur online, this study focuses on the convergence of sensor systems in the physical world. It explores the privacy implications of sensors found on our person, in our home, in our communities, and while travelling. This paper provides the following contributions: a model of human-targeted sensor systems and a framework for sensor categorization, privacy threat analysis, and countermeasure development. It concludes with a detailed case study that employs the framework to analyze the quotidian exposure encountered in an ordinary citizen's life.**

Keywords-Privacy; instrumented life; sensors; surveillance; dataveillance; uberveillance; surveillance society; panopticon

I. INTRODUCTION AND BACKGROUND

The world, and its proximate space, is becoming progressively more populated with sensors. We are aggressively, and often blindly, inviting these proliferating technologies into our environments with little understanding of how the data will be collected, shared, protected, data mined and destroyed. The reasons behind this instrumentation include efficiency, health, safety, convenience, profit, science, and security, among numerous other justifications. Tremendous value lies in the resultant data; businesses and governments are thus incentivized to collect, data mine, and retain as much as possible to accomplish immediate goals and leverage anticipated long-term objectives. Simultaneously, omnipresent sensors and perceived observers debilitate both society and individuals as exemplified by Jeremy Bentham's eighteenth century Panopticon prison design or the spiritual slavery that occurred in the Soviet Union [1,2]. Classic works on privacy and surveillance, such as George Orwell's *1984* and Aldous Huxley's *Brave New World*, have argued that governments present the greatest threat to personal freedom, but private industry and its competitive driving forces also pose malfeasance. The law frequently lags behind technological advances, and this gap allows governmental and private sector actors to maximize data collection while operating just below the threshold that invites regulatory oversight or public outcry. Nonetheless, the law is now forced to catch up. For example, the "United States vs. Jones" case recently before the U.S. Supreme Court spurred the Court to examine under what conditions police may emplace a GPS tracking device on a private citizen's personally-owned vehicle [3].

An instrumented society has many benefits, but significant detriments to both the individual and society at large exist [4,5]. Each new sensor placed in the environment, each new network link in the sensor system, and each new processing advance provide benefits to those implementing the system and sometimes to the individual and society at large, but these changes incur costs—particularly a loss of individual and collective privacy. The pervasive nature of sensors coupled with recent advances in data mining, networking, and storage technologies creates tools and data that, while serving the public good, also create a ubiquitous surveillance infrastructure ripe for misuse. Roger Clarke's concept of *dataveillance* and M.G. Michael and Katina Michael's more recent *uberveillance* serve as important milestones in awareness of the growing threat of our instrumented world [6,7].

Until recently, data was collected through manual systems and isolated analog sensors constrained by human-in-the-loop inefficiencies, but today digital sensors are scalable, networked, inexpensive, dramatically more effective, and increasingly pervasive. What were once individual islands of data now converge and aggregate into vast databases for governments, corporations, and data brokers. As this paper's concluding case study illustrates, sometimes we knowingly invite sensors into our lives to help secure our homes, track our calories, or save money, but increasingly we have little awareness that data collection takes place, a practice buttressed by policies that mandate compliance and opt-out options, all but impossible to employ or hidden in complex and legalistic privacy policies [8,9,10].

While many popular media reports examine individual classes of invasive sensors, this work takes a holistic approach by studying broadly the prevalence of sensors in our lives. We present a framework for the critical analysis, both manual and automated, of these sensor systems and for analyzing privacy-related threats. We then use this framework to conduct a case study analysis, fictionalized for privacy's sake, of common

daily activities in the context of ubiquitous sensor collection. This framework helps individuals understand the extent of monitored daily activities, often nonconsensual, and enable them to determine appropriate countermeasures. This framework helps government agencies understand the connections among sensors and better determine the effect of policies designed to protect privacy and regulate the sharing of information. It also helps information system owners understand their systems' vulnerabilities, gaps in coverage, and opportunities to leverage data by merging with other sensor networks that monitor the same subjects. While significant instrumentation and data collection occur online, academic researchers have closely studied this problem; popular media outlets have since helped inform the public and policy makers. Thus, we do not focus on online instrumentation [11,12,13]. Likewise, we do not focus on human-to-human surveillance. This paper instead focuses on quotidian exposure, via holistic analysis of electronic sensors and sensor systems of people going through their day-to-day routine in the physical world, a vital subject that has received far less scrutiny.

II. SENSOR SYSTEM MODEL

To better understand the problem, we present our model of a sensor system, as shown in Figure 1. Sensor systems collect, process, store, and disseminate information. Numerous passive and active sensors collect data about people, things, and physical phenomena, which is typically stored locally for initial processing, perhaps to do feature extraction or screening for targets of interest. The local data may be accessed by an interested (and hopefully authorized) user. If the data remains local, our concern is minimal. However, networked sensor systems transmit their data over (possibly vulnerable) communications links to remote sites for aggregation and further analysis. These remote storage sites may also be vulnerable to intrusion and exploitation. Finally, the aggregated remotely-stored data can be accessed not only by humans, but by other networked computer systems.

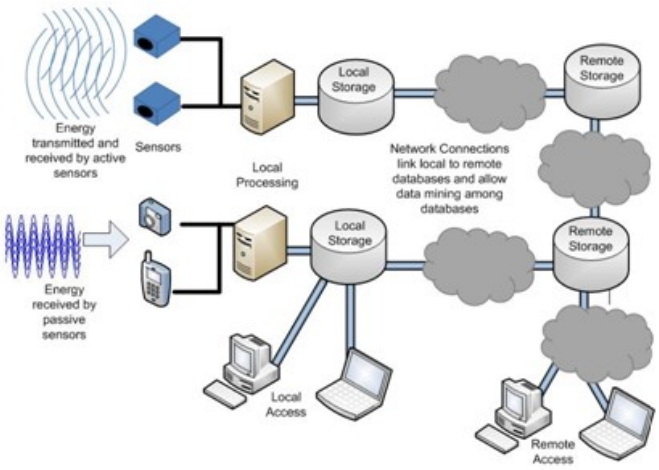

Figure 1. Overview of a generalized sensor system.

For example, a vehicle's E-Z pass tag is interrogated by an active radio-frequency (RF) sensor at the toll booth. The local processor determines an account number which it transmits to a remote processor. The remote processor queries a database system which determines whether or not the account has sufficient funds to let the vehicle pass and transmits the result back to processor at the toll booth. At the end of the billing cycle, the E-Z pass central processing system determines how much to charge each account holder's credit card and performs thousands of such transactions, interacting with their bank and credit card processing systems.

III. ANALYTIC FRAMEWORK – THE MODEL IN CONTEXT

The preceding section qualitatively describes sensor systems and their use. This sections provides a more precise but widely-applicable framework for understanding and characterizing sensor systems, including the actors who own or facilitate their use, the Information Consumers who exploit their data, the targeted Subjects, and the components and vulnerabilities of the system itself. This framework, shown in Figure 2, can be used for both generalized and focused analysis of specific instances and applications of sensor systems, as well as for the development of appropriate and effective countermeasures.

To properly construct our framework, we researched and then extended Roger Clarke's surveillance dimensions: "Of What?, For Whom?, By Whom?, Why?, How?, Where?, and When?" [14]. We found Clarke's dimensions to be a useful starting point, but insufficient to capture the full spectrum of actors, components, relationships, and attributes associated with sensor systems. Our framework contains three major parts: actors who own, enable, or consume sensor system information; the system's physical components; and its subjects. These parts are further decomposed into classes. Each class has specific attributes discussed in the following sections. There are one-to-many (1…n) relationships or many-to-many (n…n) relationships among the classes. Note as well, the self-referential notation for the sensor System class. Sensor systems often interoperate with other sensor systems. For example, a digital camera, itself a small sensor system, may be connected to a user's home computer, which aggregates data from multiple cameras, and may then share the data with other systems such as Flickr, Facebook, or Snapfish.

Although there are many ways to organize a framework, we believe actors, components, and subjects to be coherent and useful. To validate the design, we tested it under numerous scenarios ranging from very small sensor systems, such as smart phones and digital cameras, moderately complex systems such as automobiles, and large, complex systems such as the E-Z Pass automated toll collection system.

A. Actors

In the context of our framework, actors are the entities behind the manufacture, distribution, sale, regulation, and use of sensor systems, as well as consumers of the information they generate. Our framework distinguishes among four classes of actors: Owners, Enablers, Regulators, and Information Consumers.

*1) **Owner**.*

The Owner of a sensor system employs it for a variety of purposes, typically when there is an incentive, such as safety, security, or health, or a disincentive, such as regulatory fines. A single Owner can own many sensor systems, but in our

framework, we anticipate only a single Owner for each system, as indicated by the 1…n label on the line connecting the Owner class to the System class. For clarity's sake our framework shows only a single purpose for Owners, Enablers, Regulators, Information Consumers, and Subjects; but we acknowledge that purpose is multi-layered and time variant. For example, an Owner of a sensor system may employ the system for multiple reasons varying over time. Consider that an OnStar-like device may initially be employed to assist stranded motorists and generate revenue by subscription fees, but the manufacturer may later decide to sell user data to third-parties or surrender it to law enforcement or other government agencies. In these circumstances, the subscriber may not be allowed to opt-out of the information sharing or may be unaware that the company's privacy policy was quietly rewritten by corporate lawyers to allow data sharing. Thus, we see potential harm occurring frequently as purposes change. For additional insights from the legal perspective on the processing and dissemination of personal data, we recommend Solove's Taxonomy of Privacy [15]. We include the following attributes of Owners in our framework to capture their core characteristics, but these attributes may be customized or extended as desired. Listed below are representative entries for the Owner attributes:

ownerName = {City of New York, City of Los Angeles, Planet Fitness, Wal-Mart, Los Angeles Unified School District …}

purpose = {reduce shoplifting, public safety, personal fitness, ... }

role = {operator, maintainer, service provider }

The Owner effectively controls the system and sets important policy regarding its use, such as when the sensor is in operation and with whom resultant data may be shared. Note, however, that Regulators may override some policy decisions. For instance, a court of law may prescribe off-limits locations, such as clothing store changing rooms, or vulnerable or sensitive subjects that must not be surveilled, such as children and military installations. Forcible disclosure is also a possibility. For example, a night club may be the Owner of a video security system and have its video surveillance recordings subpoenaed. In addition, Owners may uninformed of all potential uses of their data and thus may be unable to make appropriate cost/benefit decisions associated with safe use of a sensor system. This issue is particularly important if the Owner of the system is also the Subject.

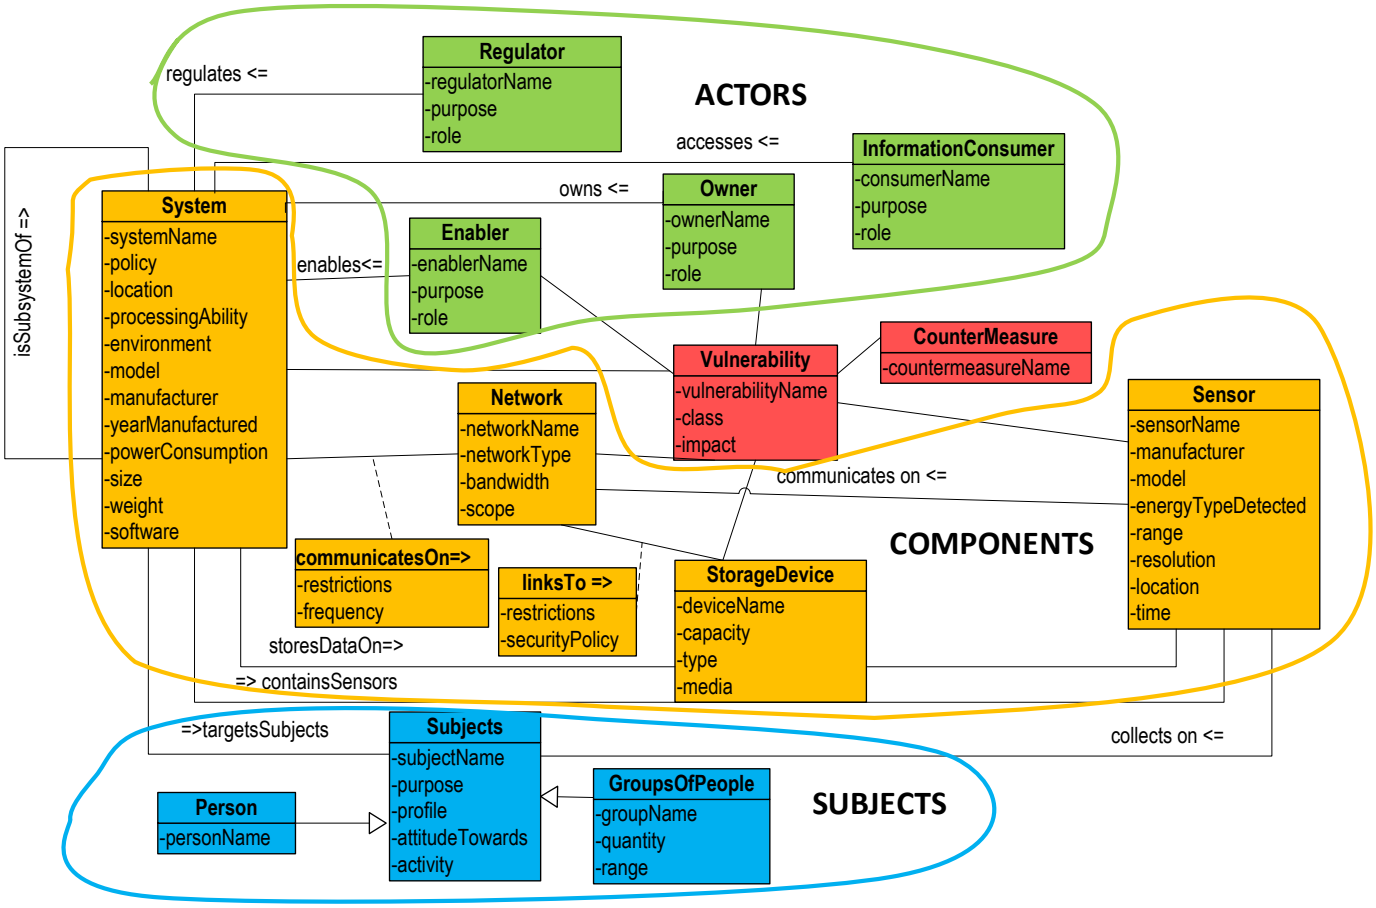

Figure 2. Sensor system analytic framework.

As depicted in Figure 2, note that the Owner, Enabler, Sensor, and several other classes have associated vulnerabilities, depicted by the Vulnerability class. This inclusion was a conscious design choice, because each of these entities possesses exploitable weaknesses that may degrade or disrupt the system's effectiveness. Likewise, for each

vulnerability there may be a corresponding countermeasure that takes advantage of the vulnerability to deny access, disrupt, degrade, or destroy the vulnerable class.

2) *Enablers*

Enablers seek to facilitate manufacture, distribution and use of sensor systems. Depending on their role, Enablers are incentivized in some way to perform these actions. For example, Nike produces and markets the Nike+ sneaker to consumers for revenue generation. Nike+ sneakers contain a small accelerometer that captures data as users run, and associated software calculates such information as calories burned, pace, distance run and elapsed time. Retail outlets sell these sneakers and Apple, in cooperation with Nike, markets a Nike+iPod sensor to enable use of specialized entertainment and performance applications. In this example, Nike, Apple, and Nike+ retail sales outlets are all Enablers of the Nike+ sensor system. Our framework supports many Enablers for each System. In our framework Enablers have the same attributes as Owners, but again these attributes may be customized or extended as desired.

3) *Information Consumers*

Information Consumers are individuals, organizations, or machines that have access to all or a portion of a sensor system's data. Whether the Information Consumer is authorized access is determined by the role attribute. Examples of roles are customer, administrator, auditor, malicious hacker (an example of an unauthorized access), or opportunist (another unauthorized access). Information Consumers utilize information for myriad purposes and possess the power, and sometimes the legal or moral right, to share information with others in ways unintended or unknown to the Subject. System policies may emplace restrictions on access or access frequency, but are not absolute guarantees against spillage or misuse.

A given sensor system may have many Information Consumers. The set of potential individual Information Consumers is very large and may include any individual, business, government, intelligence, law enforcement, non-profit organization or entity. Increasingly, we will see additional machine Information Consumers that perform some data processing or aggregation function. Human or machine aggregation is of particular concern because of the mosaic effect that occurs when multiple pieces of innocuous information from disparate sources conjoin into sensitive information.

4) *Regulators*

Regulators are organizations, often government agencies, that have a role in restricting some characteristic of the sensor system, such as what types of sensors are permissible; where the sensors may be placed; when they are allowed to collect data, who has access to the data; how, where and how long the data can be stored, and so on.

Persuading regulators to enact privacy-enhancing regulations can be a useful countermeasure to reduce the undesirable effects of networked sensor systems, as we discussed in [16].

B. *Components*

1) *Sensor*

Sensors are the heart of any sensor system and connect to the other system components in our framework. The Sensor class has several important attributes. At the most fundamental level, sensors measure energy: mechanical, electromagnetic, chemical, thermal, or nuclear. For instance, a typical camera senses visible light (electromagnetic energy). A microphone senses pressure formed by acoustic waves (mechanical energy). A household carbon monoxide detector senses chemical energy. A household smoke alarm usually detects thermal energy (heat), though it can also detect smoke using an active optical sensor. Knowing the type of energy that a sensor detects (along with other specifications such as range and resolution) allows an individual or organization to determine sensor vulnerabilities and thus devise countermeasures.

Sensors are either active or passive. Active sensors emit the type of energy that they sense and detect the energy reflected back from a target. Passive sensors have no emissions; they simply detect that form of energy already present in the environment. The distinction is important for detecting the presence of sensors and for analyzing vulnerabilities. For instance, an active sensor requires a power source and emits detectable energy. Someone searching for the presence of that sensor can simply use a passive form of that same sensor. Likewise, an active sensor is vulnerable to power disruptions. Sensors can also be classified as analog or digital. Analog sensors produce an output continuous in time and amplitude; digital ones produce a discrete-time, discrete-amplitude output. One feature of digital sensors is that the data can be further processed and aggregated with relatively little cost or effort. When this data is processed and aggregated in ways that the subject did not foresee or intend and in ways that can be harmful to the subject, this feature becomes a concern.

Two other sensor attributes in our framework are its location and time. Location is where a sensor is physically used. Location is expanding dramatically as Owners and Enablers take advantage of falling costs, the rapid rate of technological advancement, and the perceived benefit of many of these sensors. Location may be fixed or part of a mobile platform. In our framework, location includes each sensor physical placement, but it could be easily expanded to include the placement of other elements of a distributed sensor system. As we look to a future of the Internet of Things, where virtually every mundane device contains sensors and Internet connectivity, and possibly even nano-scale sensors, we could see exponential growth [17,18].

In our framework, the time attribute describes when a sensor is operational and collecting data. These time periods may be intermittent or continuous, as suggested by Clarke, triggered by a desired event such as one sensor system (like a motion sensor) activating additional sensors when movement is detected or according to a pre-programmed schedule [19]. As our case study depicts, our quotidian exposure is extensive in both time and place. Understanding when sensors operate, where they reside, and if they overlap with other sensors is necessary for analyzing vulnerabilities and developing countermeasures.

A sensor system may be composed of many diverse sensor types. Consider Apple's iPhone 4, which contains two microphones, a 5-Megapixel camera, a 3-axis gyroscope, a touch screen, and multiple antennas for Wi-Fi and cellular connectivity. The small device is both a collection of sensors and a system itself. When connected to a computer, it becomes a component in an even larger system. In summary, the attributes of the Sensor class in our framework include,

sensorName = {microphone, button, RF antenna, … }
manufacturer = {ADT, Apple, FirstAlert, Nike, Sony … }
model = {DCS932L, CO400, Nike+, ST-618 }
energyTypeDetected={mechanical, electro-magnetic, chemical, thermal, nuclear }
range = {sensor range in meters}
resolution = {in appropriate units, e.g. pixels}
location = {address or grid coordinate, …}
time = {continuous, scheduled, triggered, …}

*2) **Data and Data Storage***

Unless a sensor system is malfunctioning or inactive, sensor systems generate data, often in prodigious amounts. The StorageDevice class addresses this aspect of a sensor system in our framework.

In most cases, Owners, particularly private industry or government entities, are incentivized to collect as much new data and retain as much historical data as possible. Their intent is to mine as much information from it as possible and leave open future opportunities as data mining and processing technologies advance. In some cases, this assumption does not hold. Businesses may choose to destroy data for self-protection, as in the case of routinely destroying internal emails to protect against potential legal discovery mandates. Data may also be destroyed if an actor decides that retaining it is not cost effective, for example by applying a sliding window of data retention, if regulatory requirements require such destruction, or if some form of data anonymization is applied. However, we assert that in most cases the incentives to keep and disseminate data far outweigh any incentives to protect individual privacy by anonymization or data destruction. (And note that anonymization is often not as effective as people believe, as discussed in [20].)

Important attributes of the StorageDevice class include its name (for system access, such as c: or /home), its capacity (which may force the system to delete old data), the type of device and media used (which determines vulnerabilities). For example, the One Laptop Per Child (OLPC) project carefully considered the potential of an attacker repeatedly rewriting the device's flash memory in an attempt to destroy the drive. Example entries for StorageDevice attributes are listed below.

deviceName = {c:, d:, /home … }
capacity = {in MB or GB }
type = {Hard drive, CD-ROM, DVD-ROM, USB drive, EEPROM, SD card, …}
media = {optical, magnetic disk, tape, … }

*3) **Communication Capability and Networking***

Sensors gather data, but the data must be communicated to Information Consumers to be of value. Early analog sensors, such as a thermometer, required human observation. Today, and increasingly in the future, sensor systems include network communications: wired or wireless. These network-enabled sensor systems are easily combined into sensor networks or grids that allow vastly enhanced coverage and accuracy. Other systems, even if not directly networked, such as a digital camera, will contain large amounts of data that can be transferred manually by a digital storage technology, such as a USB drive or SD card, or automatically during a synchronization process with a larger network. The ability to communicate Sensor data is captured with the Network class in our framework. The attributes of the Network class, with exemplar entries, are

networkName = {Sprint, Verizon, … }
networkType = {wireless, USB, serial port, wired, …}
bandwidth = {10Mbps, 100Mbps, 1Gbps, …}
scope = {local, metropolitan, global, … }

*4) **System***

The System class integrates and controls the components described above. The first System attribute is its name, which could be as specific as "Jim's iPhone 3Gs" or more generic such as "Hal's home security system." Systems have policies controlling how and where the system operates; how data is stored, processed, and (possibly) destroyed; and who can access it.

The system location may be the location of the entire system or the primary control center. For example, the iPhone 3Gs is a sensor system that fits in a single enclosure and has a single location. In contrast, the E-Z Pass automated toll collection system spans several US states.

The processingAbility attribute describes how the System performs automated data analysis. In our framework, processing largely consists of data mining when data from one sensor system is aggregated with that of other systems and online databases and used by an Information Consumer. Data processing may pose the greatest risk to individual privacy. As mentioned earlier, data is frequently aggregated and rarely discarded. Data mining techniques today are extremely powerful, capable of teasing out subtleties from massive amounts of data. Whereas previous manual systems were comparatively inefficient, automated processing scales efficiently. For example, facial recognition search technology exists today, and large-scale facial recognition augmented reality technologies are expected to soon emerge [21,22]. We are also seeing crowdsourced processing systems that employ large, sometimes massive, numbers of portable sensor Systems enabled by information technology. Examples include crowdsourced crime detection and crowdsourced tagging of images on social networking sites [23,24,25]. Future technologies promise even greater gains in performance and effectiveness. Unanticipated combinations of Sensors and processing are certain. For example, researchers have programmed smart phones to use their accelerometers to monitor keystrokes on nearby keyboards with up to 80% accuracy [26]. Major contributing factors include targeted advertising and security, but automated data mining may also be used by governments, threatening individual privacy.

Processing allows identification of events, outliers, profiling, and linking disparate pieces of information to create

individual and group profiles. A key component of invasive data mining is seeking to link clusters of information to real-world identities. Such linkages may be straightforward via the use of billing records, as in the case of a user purchasing a prepaid subway card using her credit card, or it may be more difficult, requiring overt attempts at nagging the user for uniquely identifying information or the creation of social networking sites, combined with the enforcement of "true name" policies [27]. We have seen nagging data collection attempts recently on social networking and free online office suites that repeatedly request the user's telephone number. Sometimes processing performs beneficial tasks, while inviting abuse. For example, consider online dating websites that collect sufficient information such that algorithms can make effective relationship recommendations. A dating website's stockpile of highly-personal, sensitive, and potentially embarrassing information could also be used in many inappropriate ways. Examples of System attributes include,

systemName = {Jim's iPhone 3Gs, Hal's home security system, E-Z Pass NY … }
policy = {textual description of system-level policy, such as "true name" policies }
location = {address or grid coordinate }
processingAbility = {textual description of how system aggregates and processes data}
environment = {textual description of the environmental characteristics }
model = {iPhone 3GS, OnStar , Safewatch PRO RF … }
manufacturer = {ADT, Apple, IBM, Sony … }
software = {list of all software running on the system, including the operating system, with version and patch data}

C. Subjects

A sensor system uses its sensors to collect data on a subject or subjects. In our model, Subject is the human being or group under observation. Each sensor within the system seeks to measure some characteristic of the human or the human's activities, directly or indirectly by association. The energy detected by the sensor might come from an inanimate object, such as a smart card, but that inanimate object belongs to and is (presumably) located with a person. Likewise an intersection's red light sensor detects cars that have failed to stop, but each car is driven by a person. Attempts to determine the identities of individuals or to classify individuals into groups may occur during surveillance or during later processing.

In cases where a Subject is aware of a Sensor, he will frequently have a reason for tolerating, or even promoting, its existence. People buy heart-rate monitors and other sensors to increase exercise efficiency and improve health. People subscribe to services such as OnStar for safety and to increase their peace of mind when traveling. They buy cell phones and smart phones for the convenience of mobile communication. People subscribe to E-Z Pass and other automated toll collection systems for the efficiency and expediency of traveling on toll roads and bridges.

In cases where the Subject is also the Owner of a system, the Subject not only consented to monitoring, but took active steps to install and operate the system in expectation of benefits. Consent can also be subtler. For example, researchers developing facial monitoring technology suggest that Subjects may opt-in to monitoring of their facial expressions when viewing advertisements in return for product discounts [28]. If the Subject is unaware of the Sensor, Subject Purpose should be null.

When the Subject is the Owner of a system, or when he or she benefits from the system, the Subject is likely to view the system favorably. These Subjects may even support system expansion. In other cases, Subjects may be ambivalent or have no positive Purpose for tolerating the Sensor. An employee may be required to use an ID card to access a place of work. The employee may not like having his whereabouts tracked, but since that is a condition of employment, the Subject tolerates the intrusion.

Finally, Subjects may resent the sensor system or even be hostile toward it. Drivers may disapprove of radar speed traps or red light cameras. Owners of sensor systems should be aware of Subject attitudes since Subjects who hold negative attitudes toward a sensor system are likely to research and exploit vulnerabilities and attempt countermeasures. The attributes of the Subject class are

subjectName = {real world name, alias, ...}
purpose = {fitness, security, revenue, none, ...}
profile = {employee, citizen, patron, customer, ...}
attitudeTowards = {supportive, ambivalent, resentful, ...}
activity = {running, driving, speeding, stealing, ...}

D. Vulnerabilities

Each component of our system and the Owners and Enablers all have vulnerabilities. We chose to list Vulnerability as a separate class since several other classes may share vulnerabilities. For instance, most sensors, storage devices, and the system itself are vulnerable to power interruptions. Likewise, Owners, Information Consumers, Regulators, and Enablers are all susceptible to financial pressure or legal actions. Sensors that collect visible light are all vulnerable to obscuration by opaque objects (e.g. putting black tape over the lens or covering a vehicle's license plate).

Understanding the vulnerabilities of each class in our framework is necessary for critical analysis of the system. Components that are susceptible to tampering could allow unauthorized Information Consumers to access sensitive data, or alternatively allow unwilling Subjects to circumvent or disrupt the system.

Closely tied to vulnerabilities are countermeasures. Each vulnerability has one or more countermeasures that either exploit or reduce the vulnerability. We represent countermeasures with the CounterMeasure class in our framework. There is a one-to-many relationship between each Vulnerability and associated CounterMeasures. Owners, Enablers, and Subjects with vested interests in preserving the system need to understand the vulnerabilities to reduce or eliminate them. Subjects with negative attitudes toward the system or competitors interested in circumventing or disrupting

it need to understand the vulnerabilities so they can devise exploitative measures.

IV. Case Study

This case study highlights the proliferation of sensor systems throughout our society. It presents a typical day for an ordinary American citizen, Hal, a fictional, composite character used to illustrate the instrumented life. We chose to illustrate Hal to protect the privacy of the real people in our lives, but his experiences are typical of many real people.

"Ring, ring it's 6 am": Hal awakens to his preprogrammed iPhone 3GS, which he set to ring weekday mornings at precisely 0600. He rolls out of bed, checks his Facebook account on his iPhone, sets his ADT security system, and leaves for his morning workout at the local gym. His domotics-equipped, intelligent home detects Hal's departure and adjusts, records, and reports the air conditioning, lighting, and other energy-consuming home components. As he enters the gym, his membership card is scanned, recording his precise arrival time (0620), the duration of his workout, and the machines that he routinely uses. Hal connects his iPhone to his preferred Nike+iPod elliptical machine, which captures his heart rate, workout duration, calories burned, and other fitness data. His iPhone routinely sends his status to nikeplus.com, collating it with his warehoused workout profile, "keep[ing] stats on [his] every step," and tracking his "progress" and "performance" [29]. Workout complete, he updates his Facebook status and remembers to convey canned salutations to the two friends whose birthdays Facebook reminds him are today.

After departing the gym, Hal purchases a Caribbean Passion smoothie, with energy booster, at the neighboring Jamba Juice using his registered jambacard in lieu of cash. His iPhone registers his location with the cell phone towers en route, tracking his path past each tower's footprint. His E-Z Pass, acquired to avoid the lengthy toll booth lines, electronically records his passage through each toll station. Smoothie in hand, Hal fails to see the police camera positioned aside the interstate, which photographs the front and rear of his vehicle and records the date, time, and speed of his passage—later to be included on his speeding citation. The GPS feature of his iPhone tracks his progress, as does his OnStar FMV (For My Vehicle), with "Automatic Crash Response, Turn-by-Turn Navigation and Roadside Assistance," keeping Hal "safely connected while out on the road" [30].

Hal arrives at work as the security cameras capture and store images of his vehicle pulling into the parking lot and his walk to the front entrance, where his employee identification card is scanned and arrival time recorded. He enters his office, logs onto his work computer using his ID card, and begins his day of data entry—his bookmarked, multiple-tabbed, and cookie-enhanced homepages containing YouTube, Google, and Facebook open automatically. His online activities are recorded both on his work computer as well as on the remote servers that house his favorite social media sites. His work routine begins.

After a morning of data entry, tabbing to and from his diversionary websites, Hal departs for lunch, his egress captured similar to his arrival. He walks to lunch at the work cafeteria, his cell phone once again tracking his movement. Using his United Airlines MasterCard to purchase a trip to the salad bar and a Diet Coke, he joins his co-workers at their usual table. Security cameras gaze dumbly from above. Hal then returns to his office to complete his workday, textinghis wife and mother once and sending an occasional Tweet with a status update.

En route home, Hal remembers to pick up his son from the local daycare center, which logs Hal's activity for his son's protection. They stop at the local Albertson's supermarket to purchase some much needed groceries. As the strategically-placed security cameras observe from above, they purchase their items, remembering to use their Preferred Savings Card to receive their membership discount and earn bonus coupons. Hal remembers to stop at the ATM on his way out, his image and withdrawal captured by the machine's sensors.

Making their way past the E-Z Pass tolls, police traffic cameras, and cell phone towers, Hal and his son arrive home. Their domotics- and ADT-equipped home detect, record, and adjust to their arrival. As Hal prepares dinner, his son turns on their Wii gaming system, selects his Mii, and begins to play Wii Fit Plus. He notes that his cousin from across the state is logged-on and calls his father to join him. Together they join in a quick competition on the Wii's Obstacle Course, their performance shared and recorded by the gaming console. Hal's Mii reports his physical statistics and mockingly informs him that he has gained three pounds.

After Hal's wife arrives, the family eats dinner together and decides to rent a Netflix film using the Wii platform to stream the video. Movie complete, the family prepares for the night by adjusting the intelligent home's controls for the evening. Hal curls up in bed with his Kindle electronic reader and opens his recently downloaded copy of Franz Kafka's *The Trial*.

Table I provides a snapshot of some of the sensors, both self-monitored and externally-monitored, that Hal encounters during a typical workday. It lists the owner's and subject's primary purpose, along with some of the characteristics of the sensor itself (active/passive, analog/digital, type of energy detected-either electromagnetic or mechanical in these examples). The presence of these devices in our homes, at our workplaces, throughout our restaurants, shopping centers and recreational facilities, and alongside our highways and sidewalks continues to proliferate. As these sensors grow in their physical presence and sophistication in gathering, warehousing, and collating data about our quotidian existence, our digital profile grows increasingly revealing and, correspondingly, increasingly profitable for commercial and government entities to exploit. As the case study reveals, many of these sensors overlap in both purpose (e.g., gym membership and Nikeplus.com) and in area surveyed (e.g., E-Z Pass, police cameras, cell phone towers), with little of Hal's activities or movements left unmonitored by a device emplaced by him or another actor, known or unknown. Many of these devices exist either without Hal's knowledge or overlooked by him due to their familiar presence during his daily travels. He remains blithely unaware of their constant digital mappings and his maturing digital profile.

TABLE I. SENSOR SYSTEMS ENCOUNTERED BY SUBJECT DURING CASE STUDY.

	Owner Purpose			Subject Purpose		Sensor Characteristics		
Sensor System	Track Movemt	Health	Habits, Ads	Security	Convenience	Active	Digital	Energy Detected
iPhone 3GS	x	x	x	x	x	x	x	emag, mech
ADT Security System	x		x	x		x	x	emag, mech
Gym Membership Control system	x	x	x				x	emag, mech
Nike + IPod and Nikeplus.com	x	x	x					emag, mech
Jamba Juice jambacard		x	x		x			emag
Cell Phone System	x		x		x			emag
EZ Pass	x		x		x		x	emag
Law Enforcement Cameras	x	x	x				x	emag
GPS receiver	x		x				x	emag
OnStar	x		x			x	x	emag
Work Place Security Cameras	x		x			x	x	emag
Work ID Card Scanner	x		x		x		x	emag
Computer Logon Recorder	x		x				x	mech
YouTube			x		x			emag
Google			x		x			emag
United Airlines MasterCard	x		x		x		x	emag
Day Care Security Sys.	x		x				x	emag, mech
Albertson's Preferred Savings	x	x	x		x			emag
ATM Camera and Log	x		x			x	x	emag, mech
Wii Nintendo		x	x		x	x	x	emag, mech
Kindle (with microphone)	x		x		x		x	emag, mech

V. COUNTERMEASURES

The question remains: How might Hal counter the proliferation of these sensor systems that record his nearly every activity? The search for potential countermeasures begins by examining the vulnerabilities of each of the framework's components. For example Sensors each possess unique characteristics that can be defeated through a variety of techniques, such as shielding, detection, spoofing, jamming, even physical destruction. Subjects frustrated with the presence of sensor system systems may employ these techniques to reduce a sensor's or system's effectiveness. Not all countermeasures are technical or even physical; some seek to alter the incentives of Owners and Enablers. Subjects, or other interested parties, may boycott businesses or target political support to change surveillance practices. Well-conducted research can effectively challenge an Owner's purpose, as occurred in Los Angeles when statistics showed that financial and safety benefits were not being achieved [31]. Even public shaming has proven to be effective [32]. Subjects may even band together and collaboratively track locations of Sensors, such as the New York City Surveillance Camera Project, which publishes a detailed map depicting camera locations [33,34]. Countermeasures are discussed in more detail in [16].

VI. EXTENDING THE FRAMEWORK

By necessity, our framework is easily extensible. The rapid rate of technological change and increased adoption of sensing technologies by individuals, communities, governments, and businesses demand a framework that is adaptable.

Similarly, due to space constraints, we did not include exhaustive listings of possible values for each attribute. Nonetheless, we attempted to incorporate comprehensive sets for attributes that contain a relatively small number of potential values. For other attributes that might have hundreds, thousands, or more potential values, we attempted to provide diverse representative examples that, again, could be easily augmented as necessary.

VII. CONCLUSION

In the face of sampling of the physical world by ubiquitous sensors combined with online tracking, targeted advertising, well-intentioned (and less well-intentioned) governments, law enforcement agencies, and companies, increasingly smaller portions of our lives remain private and unmonitored. This paper analyzed the key actors, component parts, and subjects of sensor systems along with important attributes, purposes, and vulnerabilities from a variety of important perspectives. These insights enable current and future critical analysis of sensor systems, which can lead to more effective countermeasures, metrics for quantifying and articulating sensor spread and quotidian exposure in the day-to-day world, and a better understanding of incentives, benefits, and harms from an instrumented world. We acknowledge that the examples we provided in this paper will be replaced by new appearances of sensors in our environment. The framework is designed to flexibly accommodate new developments.

When considering the societal impact of an instrumented world, Robert Frost's "The Road Not Taken" comes to mind: "I shall be telling this with a sigh. \ Somewhere ages and ages hence: \ Two roads diverged in a wood, and I-- \ I took the one less traveled by, \ and that has made all the difference." In an

instrumented world, this idealistic view is replaced with a world of sensors and Markov models that determine the probabilities of all travelers on all forest paths to determine where to best place advertisements or to determine the best sensor for a security checkpoint.

ACKNOWLEDGMENT

The views in this article are the authors' and do not necessarily reflect the official policy or position of their employers.

REFERENCES

[1] Bentham, Jeremy. The Panopticon Writings. Ed. Miran Bozovic, London, Verso, 1995. http://cartome.org/panopticon2.htm (accessed 15 Nov 2011).

[2] Aron, Leon. "Everything You Think You Know About the Collapse of the Soviet Union Is Wrong." Foreign Policy, July/August 2011.

[3] Liptak, Adam, "Court Casts a Wary Eye on Tracking by GPS," New York Times, 8 November 2011.

[4] Shilton, Katie. "Four Billion Little Brothers." ACM Queue, August 2009, Vol. 7, No. 7.

[5] Solove, Daniel. "I've Got Nothing To Hide and Other Misunderstandings of Privacy." San Diego Law Review, July 2007, Vol. 44, p745.

[6] Clarke, Roger. "Information Technology and Dataveillance." Communications of the ACM, May 1988, Vol. 31, No. 5, pp. 498-512.

[7] Michael, M.G. and Katina Michael. "Towards a State of Uberveillance." IEEE Technology and Society, Summer 2010, Vol. 29, No. 2, pp. 9-16.

[8] Popken, Ben. "Giant List of Data Brokers to Opt Out Of." The Consumerist, 21 June 2010.

[9] Catone, Josh. "Trying to Decipher that EULA? Better Have a PhD." Sitepoint, 4 September 2008. http://www.sitepoint.com/trying-to-decipher-that-eula-better-have-a-phd/ (accessed 14 July 2011).

[10] Doctorow, Cory. "Video-game shoppers surrender their immortal souls." BoingBoing, 16 April 2010. http://www.boingboing.net/2010/04/16/video-game-shoppers.html (accessed 15 July 2011).

[11] "What They Know." Wall Street Journal, Investigative Series, 2010-2011. http://online.wsj.com/public/page/what-they-know-digital-privacy.html (accessed 4 Oct 2011).

[12] Krishnamurthy, Balachander and Craig Wills. "Privacy Diffusion on the Web: A Longitudinal Perspective." International World Wide Web Conference, April 2009.

[13] Conti, Greg . Googling Security. Addison-Wesley, 2008.

[14] Clarke, Roger. "What is Uberveillance (And What Should Be Done About It?)." IEEE Technology and Society, Summer 2010, Vol. 29, No. 2, pp. 17-23.

[15] Solove, Daniel. "A Taxonomy of Privacy." University of Pennsylvania Law Review, January 2006, Vol. 154, No. 3, pp 477-560.

[16] Shay, Lisa and Greg Conti. "Countermeasures: Proactive Self-Defense Against Ubiquitous Surveillance." HOPE 9, New York City, July 2012 http://www.rumint.org/gregconti/publications/countermeasures_v89a.pdf (accessed 17 July 2012).

[17] Ashton, Kevin. "That 'Internet of Things' Thing." RFID Journal, 22 June 2009.

[18] Kurzweil, Ray. The Singularity Is Near: When Humans Transcend Biology. Penguin, 2006.

[19] Clarke, Roger. "Information Technology and Dataveillance." Communications of the ACM, May 1988, Vol. 31, No. 5, pp. 498-512.

[20] Ohm, Paul, Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization (August 13, 2009). UCLA Law Review, Vol. 57, p. 1701, 2010; U of Colorado Law Legal Studies Research Paper No. 9-12. http://ssrn.com/abstract=1450006 (accessed 17 July 2012)

[21] "Anonymous No More." The Economist, 30 Jul 2011.

[22] Eaton, Kit. "The New iPhone's Facial Recogition Capabilities Could Redefine Privacy." Fast Company, 4 October 2011. http://www.fastcompany.com/1785016/face-recognition-apples-secret-weapon-to-change-social-media-gaming-online-privacy (accessed 19 October 2011).

[23] von Ahn, Luis. "Games with a Purpose." IEEE Computer, June 2006, Vol. 39, No. 6, pp. 96-98.

[24] Paul, Ian. "Facebook Photo Tagging: A Privacy Guide." PC World, 9 June 2009. http://www.pcworld.com/article/229870/facebook_photo_tagging_a_privacy_guide.html (accessed 19 October 2011).

[25] Internet Eyes, CCTV Monitoring Home Page. http://interneteyes.co.uk/ (accessed 10 November 2011).

[26] "Georgia Tech Turns iPhone into spiPhone." Press Release, Georgia Institute of Technology, 18 October 2011. http://www.gatech.edu/newsroom/release.html?nid=71506 (accessed 15 Nov 2011).

[27] Blue, Violet. "Google Plus Deleting Accounts En Masse: No Clear Answers." Pulp Tech Blog, ZDNet, 23 July 2011. http://www.zdnet.com/blog/violetblue/google-plus-deleting-accounts-en-masse-no-clear-answers/567 (accessed 19 October 2011)

[28] "The All-Telling Eye." The Economist, 22 October 2011,

[29] Nike + iPod. Product web page, http://www.apple.com/ipodtouch/built-in-apps/nike.html (accessed 19 October 2011).

[30] OnStar For My Vehicle product website: http://www.onstar.com/web/fmv/home (accessed 16 Nov 2011).

[31] Rubin, Joel. L.A. Traffic Cameras May Get the Red Light. Los Angeles Times, 8 June 2011, http://articles.latimes.com/2011/jun/08/local/la-me-0608-red-light-20110607 (accessed 14 July 2011).

[32] Musil, Steven. "Lawsuit accuses Cisco of aiding Chinese repression." CNET, 22 May 2011. http://news.cnet.com/8301-1023_3-20065219-93.html (accessed 15 Nov 2011).

[33] New York City Surveillance Camera Project. http://www.mediaeater.com/cameras/ (accessed 22 October 2011).

[34] Note that some police departments publish the time and location of DUI checkpoints in an attempt to deter drunk driving.

LISA A. SHAY is an Assistant Professor in the Department of Electrical Engineering and Computer Science at the US Military Academy at West Point. She holds a B.Sc. from the US Military Academy, an M.Sc. from Cambridge University, and a Ph.D. from Rensselaer Polytechnic Institute, all in Electrical Engineering. She is a Senior Member of the Institute of Electrical and Electronic Engineers and a licensed professional engineer.

GREGORY CONTI is an Associate Professor in the Department of Electrical Engineering and Computer Science at the US Military Academy at West Point. He holds a B.S. from the US Military Academy, an M.S. from Johns Hopkins University, and a Ph.D. from the Georgia Institute of Technology, all in Computer Science. He is a Senior Member of the Association for Computing Machinery.

JOHN NELSON is an Assistant Professor in the Department of English and Philosophy at the US Military Academy at West Point. He holds a B.S. from the US Military Academy, a M.A. from Oregon State University, and a Ph.D. in Comparative Literature from University of Washington.

DOMINIC LARKIN is an Assistant Professor in the Department of Electrical Engineering and Computer Science at the US Military Academy at West Point. He holds a B.S. from Troy State University and a M.S. from the Georgia Institute of Technology, both in Computer Science.