

TIPPERS: A Privacy Cognizant IoT Environment

Sharad Mehrotra, Alfred Kobsa, Nalini Venkatasubramanian

Donald Bren School of Information & Computer Sciences
University of California, Irvine, CA, U.S.A.
{smehrotr,kobsa,nalini}@uci.edu

Siva Raj Rajagopalan

Honeywell Automation and Control Solutions
1985 Douglas Dr. N.
Golden Valley, MN, U.S.A.
siva.rajagopalan@honeywell.com

Abstract—IoT environments are important and challenging application domains for privacy studies, for two reasons: data collection about people is nearly invisible, and IoT environments typically do not support interfaces that allow users to specify their privacy preferences or to control the personal data collection practices of the environment. It is also well known that IoT environments present rich challenges in managing privacy choices. We present two systems that will support powerful mechanisms to embed and test a diverse set of privacy technologies in an IoT environment in a large School building.

Keywords—IoT environment; privacy

I. INTRODUCTION

The objective of the TIPPERS project is to provide two test beds for privacy research in the Brandeis [1] program: an existing system and a research system that support powerful mechanisms to embed and test a diverse set of privacy technologies. The application domain of both systems is the Internet of Things (IoT). IoT environments are very important and challenging application domains for privacy studies, for two reasons: in IoT environments, data collection about people is nearly invisible, and, furthermore, IoT environments typically do not support interfaces that allow users to specify their privacy preferences or to control the personal data collection practices of the environment. It is well known that IoT environments present rich challenges in managing privacy choices [2].

Regarding the **research system**, we plan to install a large number of off-the-shelf environmental sensors for presence, location, identity, and event and activity recognition in Bren Hall, a six-story office building on the UC Irvine campus. The research system will also include the wearable devices of all occupants of the building. The targeted building houses various types of occupants: administrative staff, technical staff, professors, graduate students, visitors-in-residence from industry and academia, and many who are in the building for a very short time only. Potential services of such an IoT environment would include: an evacuation tally count in the case of fire and earthquakes; a current-location directory as well as automated proximity alerts to facilitate face-to-face encounters; thermostats that learn and predict people's presence in their offices and regulate the room temperature accordingly; usage heating, ventilating, and air conditioning; attendance recording for students in classes and/or seminars; tracking potentially suspicious intruders or activity reported on campus; analysis to understand social dynamics of building and its occupants, etc. All these potential services are obviously somewhat privacy-intrusive. The research system will be designed and developed

using Honeywell's Tridium open IoT platform that is a commercial leader in the IoT device integration area, and SAT-WARE [3], a semantic middleware for sensor data processing. The research system will leverage these technologies to build a fully-functional data acquisition, management and analysis framework and incorporate mechanisms to intercept/route dataflow through various privacy technologies.

Regarding the **existing system**, it will use the same sensor base and physical location as for the research system, but would use the Honeywell Enterprise Building Integrator (EBI) [4]. EBI is a management interface for a variety of subsystems in a building such as heating, ventilation, and air-conditioning (HVAC), lighting, fire and life safety, occupancy monitoring, access control and surveillance video. EBI is designed as an integrator and can interface with a variety of third-party sensors and actuators and management systems; it allows access to its internal data via easy to use interfaces. The proprietary EBI based solution will be installed as a commercial product so that we have an actual real-life existing system that is currently managing thousands of buildings across the world. This system will be retrofitted with mechanisms to support privacy technologies and privacy research. First, data from all managed building sensors can be accessed easily at a central location enabling joint analysis of multiple data streams to find new opportunities for privacy models and mechanisms. Second, EBI allows application access to the stored data, so that the results of privacy mechanisms on the applications that use the data can be viewed side by side with the unaltered system.

II. RESEARCH APPROACH

Unprecedented growth in sensing, data capture devices, communication and computing technologies has created a possibility that in the near future we will be able to continuously capture and analyze (in real-time) almost every aspect of our lives: personal experiences, social interactions, and our interactions with engineered, cyber, or physical systems or the environment. While interconnected sensors and devices embedded in the environment, wearable technologies, social networks, and data generated from human-machine interactions (e.g., click stream data and audit logs) create limitless possibilities, one of the key concerns/challenges in developing an information infrastructure for live data applications is the loss of privacy and confidentiality. Technologies in such a context must therefore address three fundamental tasks:

1) analysis and understanding of vulnerabilities and inference channels that can lead to risk of privacy breach,

2) development of privacy protecting technologies to hide sensitive information while still enabling the end-goal for which data is being shared, and

3) regulations and policies that promote privacy and security of sensitive data (where users have reasonable expectation of privacy).

Though there is increasing awareness of privacy risks and protection mechanisms for sensory data, current solutions are only applicable to traditional data sharing applications (e.g., sharing medical (patient) data and census data to support data analyses). Sensor-rich spaces offer a new set of challenges. The semantic richness of pervasive information-rich sensors introduces new inference channels (e.g., see [5]). Hidden in such concerns are architectural issues of where and how sensor data processing is performed. For instance, if data in its raw form is not externalized and only inferences are shared, the problem is significantly simplified. However, externalizing raw sensory data may be necessary for a variety of reasons, e.g. the application of proprietary inference algorithms. For example, Apple’s Siri intelligent personal assistant and Google’s Voice Actions requires raw sound signals to be sent to cloud servers for speech recognition. Such data can, however, enable inferences about identity, age, gender, location, emotional state, etc. Likewise, video data shared might allow for inferences about personal habits, associated objects, clothing choices, gestures, many of which might be deemed sensitive. Additional challenges arise due to the continuous nature of sensor data, and from the ability to fuse multiple diverse sensory data.

Our goal is to develop a comprehensive framework and model under which privacy risks from sensory data can be studied and corresponding protection mechanisms (technologies to support privacy-utility tradeoffs) can be designed in a concrete systems context. We explore and compare two related systems:

1) an existing IoT data processing system based on Honeywell EBI technology that provides state of the art data collection and analysis in a real-life system that can be easily accessed and modified for privacy research, and

2) a research IoT data processing prototype that supports flexible mechanisms and APIs to intercept data flows amongst diverse software and hardware modules to embed a variety of privacy technologies thereby realizing the goal of designing a system that supports privacy-by-design principles.

The next sections elaborate on the system and the test bed.

III. THE EXISTING SYSTEM

The existing system will be based on the core of the Honeywell EBI system, which will be installed in the same building as the research system but provides the perspective and capabilities of a state of the art building management system that is currently used in thousands of buildings worldwide. This system will collect sensor data from possibly hundreds of sensors across this large building that contain information about the activities of its inhabitants. This commercial system will be extended to enable other researchers to experiment with and validate various privacy mechanisms and metrics so that it can serve as a unique open platform for IoT privacy research.

A. The Honeywell EBI System

Honeywell Enterprise Buildings Integrator (EBI) [4] can be comprised of one or more of the following applications on a single server or a network architecture of servers (see Fig. 1): Building Manager provides HVAC control, Security Manager provides interfaces to Access Controllers, Life Safety Manager manages Fire Alarm Systems and Smoke Control, and Digital Video Manager manages all video and audio sensors in both real time and archival modes. EBI provides a consolidated view of all these systems and acts as a single point of control for quick response. EBI can be used to conceive and script complex system relationships both for operational efficiency and for emergency response. Thus, the security system can maintain overall site protection, while unlocking specific doors based on the alarm location to speed egress and allow quick access by emergency responders. Digital video and audio can provide a comprehensive assessment of the alarm location as soon as the alarm is received. These examples show the immense value of an integrated management system in a building or campus.

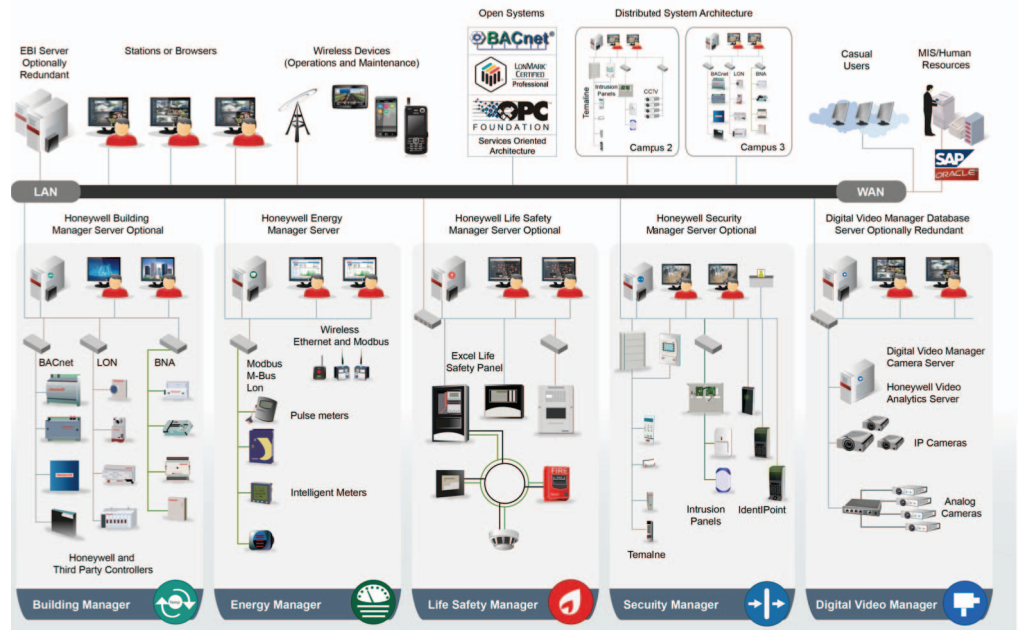


Figure 1: Honeywell Enterprise Building Integrator

B. Retrofitting the existing system with privacy technologies

EBI provides a range of features and interfaces that enable the creation of “plug-in” mechanisms of any kind, which can be used for our purposes of retrofitting a range of privacy technologies. EBI provides open database connectivity (ODBC) access to its database, so that any program that wants to analyze collected sensor data can access it through this interface. Extending the functionality of EBI for new capabilities can be achieved in a variety of ways. EBI provides a suite of tools, which enable the customers to create their own powerful applications on top of EBI’s functionality. The EBI architecture will allow a diverse set of privacy mechanisms to be incorporated to support policies and preferences of individuals in the immersed space. For instance, it will enable us to obfuscate/hide data possibly through randomization and/or noise addition to implement privacy guarantees. The logs of activities generated will further enable us to determine potential privacy inferences, which can then be used by the privacy broker to implement privacy policies.

IV. DESIGN APPROACH FOR THE RESEARCH SYSTEM

Privacy concerns in IoT systems arise when users do not fully trust the organizations that provide the IoT service, the humans who operate the environment, or users or applications who may have access to data and higher-level information. The implied dangers of not addressing privacy concerns in IoT applications make a compelling argument for privacy mechanisms to be incorporated as an integral part of IoT system design. In the context of data sharing applications, by privacy we refer to limiting or preventing disclosure of attributes or information about individuals that is deemed as sensitive. However, the challenge is not just to protect data that refers to attributes addressed in privacy policies, but also to establish whether an adversary can infer sensitive knowledge from pieces of information that are by themselves not sensitive.

Consider the privacy policy of an individual Bob who does not wish Alice to know when (and how often) he visits the smoking-lounge in the office building. Naturally, Bob’s policy will disallow Alice to get an update from the sensor at the entry/exit of the smoking-lounge. Information about Bob’s presence in the corridor leading to the smoking room is likely not sensitive in this privacy policy and, thus, could be revealed. However, if this corridor leads to the smoking-lounge only and not to any other rooms, knowledge of Bob’s presence in the corridor would reveal his visits to the smoking room.

Disclosure of sensitive data or sensitive inferences can occur when data flows from one device/software component/user to another and, in turn, depends upon the form in which such data is shared, as well as pattern of access to data. Mechanisms to intercept data/access requests between components can serve as a powerful test bed for studying privacy technologies in IoT settings. Our research system explores a “data management centric” view towards the design of such IoT architecture. We identify key data management needs and limitations of existing technology and propose a novel system architecture aimed at overcoming the limitations. The design of the next generation research platform aims to reduce any inbuilt legacy requirements and therefore naturally provides a mechanism for clean-slate design of privacy tools and metrics. It will aim to support opportunities for multiple types of data privacy (e.g., large data,

small data, multiple participants, etc.) and present rich challenges in managing privacy choices.

From a data management perspective, IoT applications follow a standard data processing pipeline where raw data (collected from diverse types of sensors) is collected, analyzed, transformed, stored, and then consumed by applications. Such data processing pipelines can range from real-time (e.g., in the context of dynamic observation and control) to offline processes that allow systems to collect and create databases that can then be analyzed for variety of purposes. While data processing pipelines in IoT applications are structurally similar to those supported by modern data management and analysis systems, the nature of IoT data and applications offers significant challenges and also opportunities for new innovations in system design. The proposed research system is driven by three fundamental observations:

1) Traditional information system architectures, where data is first collected into large data warehouses and then made available for analysis and awareness, will not meet the real-time and dynamic control needs of a large number of live-data applications. Furthermore, such a collect & build approach requires trust in the warehouse (or, alternatively, pushes designs to only consider encrypted data processing at the Warehouse).

2) Given the complex pipelines needed for real-time analysis of live data, it may be impossible to capture and process live data in its entirety. The execution of pipelines must adapt to the ever-changing needs of their applications as well as the characteristics of the data produced. Thus, the system has to support to address veracity of the data.

3) Sensor data processing (specially in the context of semantic sensors as cameras) incurs errors in analysis and event detection that potentially impacts application quality. Such errors could be a manifestation of incorrect/incomplete data capture, and/or be due to limitation of data interpretation and information extraction techniques. One approach to model such uncertainty and application impact is through modeling trust in observations and events.

Mechanisms to address the above observations pose unique challenges and opportunities vis-a-vis privacy. The structured pipelines for IoT data processing offer clean and logical intercepts for the insertion of privacy solutions. For example, processing may be performed locally at the sensor (e.g., using trusted local processing as in a cell phone) and such a solution may offer improved performance and privacy without necessitating trust in the database server. Similarly, inbuilt approaches to deal with data uncertainty may be leveraged to incorporate privacy (that may introduce uncertainty) in the data path. Techniques to enforce privacy must be carefully chosen though. E.g., data may be further suppressed due to privacy constraints, and this can further reduce trust/veracity of the data/events and thus result in lower application utility.

Our research studies a novel approach to modeling and representing live sensor data to improve system performance and control privacy leakage. We will design an end-to-end system architecture in which data producers (viz. devices/sensors) are considered a unified part of the overall system, and data processing techniques take an end-to-end view on optimizing and controlling the full data pipeline for both utility and privacy.

B. Privacy by Design Interventions in the research system

In the research system, data (possibly carrying personalized information about subjects) flows from sensors through diverse extraction and analysis software where it is aggregated and analyzed for semantically higher level events, to individuals and applications who request the data for an application-specific purpose. The hardware and software platforms through which the data passes through and is aggregated/analyzed might be at different levels of trust – e.g., the database of events or the logs may run at trusted servers or may be located over the untrusted cloud. Likewise, subjects whose identifying information migrates across myriads of devices, software, applications, and users may have information-sharing policies that restrict what information is revealed about them. Further, end-users/applications may also have privacy requirements to hide their information needs from the system. Our approach aims to restructure our relationship with data by shifting the mechanisms for data protection to the data owner rather than the data user".

To empower the envisioned system to serve as a vehicle for a diverse array of privacy technologies and solutions, the research system will support powerful mechanisms and APIs to intercept data and information flow amongst components of the system and to embed a variety of privacy interventions. The intercepts can be used to transform what data, in which context, and in what form migrates across diverse hardware/software boundaries. For instance, in an application context such as understanding the occupancy levels of different parts of the building, where an application may ask for location data about individuals from the database, the resulting answers may be intercepted and passed through differential privacy mechanisms to remove identifying information. As another example, consider policy specification and enforcement to limit disclosure of sensitive information to others in the context of data sharing. Our research system, with the support for multiple levels of information abstraction, offers opportunities to explore how policies can be expressed at different levels (e.g., at the event and entity level) and how such policies can be enforced when data is queried at a lower abstraction level such as at the sensor level (e.g., "give me the stream of data generated by sensor S1"). The research system not only provides mechanisms to control the data flow between components to implement privacy policies, but it also offers diverse levels of information abstraction to specify and reason about inferences that may result due to data sharing. The research system, due to its flexibility, empowers us to explore the related privacy challenge using a widely different privacy technology. Consider again, the challenge of privacy preserving data sharing in the context of sensor (and/or event data). In contrast to a policy based approach where the goal is to maximize information being shared while enforcing the policies, an alternate technology is that of "minimum disclosure" wherein the objective is to provide the least information to adversaries that suffices for the task of the individuals while at the same time minimizing information disclosure. Such minimum exposure techniques have been explored in the literature including by us as in the case of surveillance applications where we were able to upper-bound the "extra" information disclosed to the adversary in the context of privacy preserving event detection. Such techniques, as well as extensions, that facilitate a tradeoff between the amount of information disclosure and the probability with which an

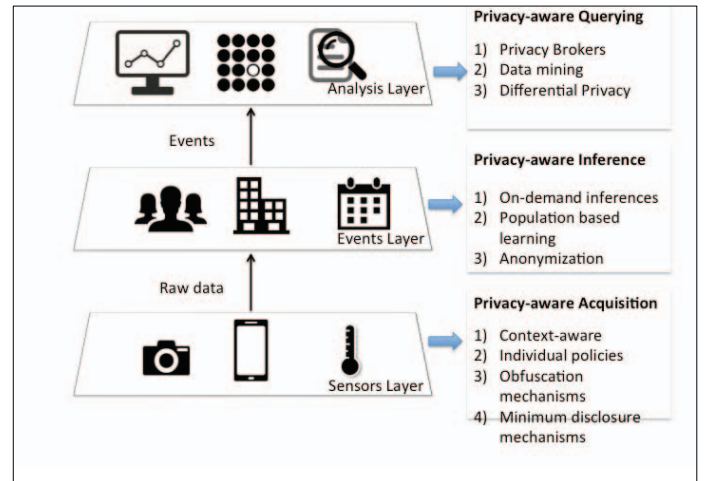


Figure 2: Levels of abstraction supporting privacy

adversary or untrusted user can determine if an event actually occurred (viz. utility) can be easily incorporated into the proposed research system.

Fig. 2 further highlights the diverse levels of data abstractions supported in the research system, and some of the privacy challenges that such a data representation allows to explore. In addition, the research system will support the integration and study of encrypted database search (e.g., encrypted log search) and secure multi-party computation.

While one can design the ability to insert privacy techniques at different levels of the system stack, the key challenge lies in ensuring the choice of strategies and their joint execution in the end-to-end flow supports both the user privacy and application utility goals of the IoT deployment. In the following sections, we discuss such an end-to-end design of a privacy cognizant IoT research system and its implementation.

C. Designing IoT Data Management to support plug and play privacy techniques

The proposed IoT data management system takes an end-to-end view of the system. From an architectural perspective, the system consists of devices from which data may flow, to processing units where the data is analyzed (e.g., to enrich the data, to clean it, to merge it with other data sets, etc.), to storage system where the data may reside, to IoT applications that consume the data. The first question we face is what kind of data model should such a system support.

One approach is a data model/abstraction that allows a certain level of separation of concerns. Seen from the IoT application, the various devices and sensors are simply data capture devices. Such devices are not intrinsic components of the application logic. For instance, an application monitoring location of a person is interested in the location and not in the specific sensor used to monitor the location. So one can imagine modeling the physical world/domain, in the same way we model domains in current databases, as physical entities and relationships. The difference is that we are now modeling the dynamically evolving physical world. To capture the dynamicity, we differentiate between immutable attributes and attributes that are dynamic and change as a function of time. For instance, a

person's name is immutable, but his/her location changes. Likewise, relationships between entities may be dynamic – e.g., if the system captures the fact of a person entering a room, then based on the movement of a person a new relationship between a person and a room may dynamically emerge. Now imagine the dynamically evolving world represented in such an extended database which understands the notion of data evolution. Applications can almost entirely be built on such an abstraction, without having to deal with the specificity of how the dynamic data was captured. Regarding applications, the data might have been fed by sensors, may have been explicitly entered by a human, or could have been a result of a simulation of a prediction model. Indeed, this framework provides a natural way to exploit predictability of the natural phenomena into data processing.

Associated with the underlying data (specifically its dynamic properties) are “sensors” which can be tapped into to observe the “value” of the attribute/relationships. Different sensors may differ in the quality of the observations they produce and may differ in the cost of generating the observation. So a natural question in such a setting would be to compile the application logic (based on the higher level semantic data representation) and from that derive a sensor data capture plan. This decoupled representation of sensors and application provides us with a natural mechanism to optimize sensor data acquisition, similar to techniques studied in the past such as multi-query optimization. Such a plan must consider the desired quality needs/tolerances of different applications.

In the above approach, the system must maintain an explicit representation of sensors currently available, including the coverage (i.e., what can they sense) as a function of time. The system maps the applications needs for dynamic attributes/relationships of the physical world to the sensors whose coverage can support the desired quality/need of the application. Two things are further important about the above model –

1) The way described above, the system models data into two layers – a semantic layer which models the physical world and its dynamic phenomena (represented as dynamic attributes and relationships) and the lower level sensors which produce streams to observe the dynamic phenomena (such sensors may dynamically join/leave). The system exploits quality, coverage, needs, application tolerances, etc. to generate a schedule for data capture. There is no real need to limit semantic abstraction to merely two layers. One can imagine a layered approach that generates different abstraction levels suitable for different types of applications and the data being transformed across layers in a principled way. In other words, such a model can support concepts such as rollup, drill down, etc. common in OLAP types of applications.

2) The data model that decouples semantic concepts (and hence application logic) from sensors/devices was primarily motivated to overcome the burden of dealing with sensor/device heterogeneity from application programming (since in the envisioned model, applications specify their data/information needs and the system, which maintains device information, scope, coverage etc. maps the need to appropriate data capture plans), and scale (the data capture rates/actuation parameters are set based on applications needs dynamically). An additional advantage of the proposed model is that it provides a natural framework to implementing privacy policies in sensor based

systems. Privacy policies and inferences are a lot easier to understand, specify, and reason with at the semantic level when the underlying data is interpreted into semantically meaningful observations in contrast to the (uninterpreted) sensor level. We can envision an enhanced model that allows specification of policies (as well as learning of inferences) at the semantic layer which are then translated into making decisions on whether sensor data or its interpretation should or should not be shared across different trust boundaries.

D. Implementing the research system

To explore tradeoffs amongst privacy, trust and information needs, and to explore the efficacy and validity of our research at an experimental level, we will build a system that supports association of privacy policies for individuals immersed in the environment and trust in events generated. The system will leverage the existing SATware system, which serves as a semantic middleware for sensor data processing, and Tridium, Honeywell's open IoT platform. We briefly describe these systems next. The SATware System [3] is a scalable pervasive space middleware, which provides seamless access to sensor and event level data. Applications access this information via a SQL style query language referred to as SATQL, at both the physical (e.g., raw sensor feeds) and semantic levels (i.e., at the level of entities, activities, and events). The key concept is that of a virtual sensor that empowers programmers to define and detect semantic concepts, thereby realizing information abstraction as discussed in Section IV. Virtual sensors are mapped at run-time to a graph of operators which are implemented over physical sensor streams. SATware, in addition, contains a SATDeployer component that “optimally” deploys operators to various nodes in the system to meet application quality requirements.

The Tridium JACE controller with the Niagara Framework is a unique Java-based platform that allows one to develop custom applications from scratch for accessing, automating and controlling “smart” devices in real-time over a network if needed. By converting connected building system sensors and actuators (and their data and attributes) into software objects regardless of make, model and manufacturer, it is possible to read real-time data, send commands to the device and utilize common programming tools to reconfigure and reprogram them. While the preferred hardware platform is a Tridium JACE controller, the Niagara Framework can operate on a common PC so that experiments can be conducted on virtual devices. This would be especially useful for conducting experiments to validate privacy mechanisms on simulated data before trying them in live experiments. By integrating protocols that are peculiar to building controllers into the stack, it frees the application developer to focus on higher-level functions and abstractions while having access to information from heterogeneous devices, legacy systems, etc. Niagara provides powerful and unrestricted access to all information that can be gleaned from sense and control devices and then mined for use by any software. JACE and Niagara together thus allow privacy researchers to create their own sensors with privacy enhancements and applications that are privacy-aware, and have these run side by side with existing sensors and applications that are privacy-challenged. More details can be found at www.tridium.com.

Using the SATWARE middleware and the Tridium platform as starting points, we will build a prototype pervasive space platform that provides open interfaces to embed a variety of privacy and trust technologies leading to the envisioned research system prototype. The system, along with application and demonstration studies built using it, will serve as a catalyst for understanding privacy challenges, and provide a concrete context to incorporate a variety of privacy technologies and to test them in isolation as well as in the integrated study. The system will also serve as a vehicle to develop privacy metrics and study the efficacy of such metrics in capturing privacy perception of users through dedicated user studies.

V. TESTBED AND APPLICATION DESCRIPTION

Both the research and existing systems for IoT data management and the application development will be deployed in Bren Hall at UCI, a 90,000+ square feet 6-story building that houses UCI's School of Information and Computer Sciences. Already, we have a significant sensor deployment in the building (see <http://www.i-sensorium.org>), including video cameras, sensor mounted mobile robots, people counters, RFID, acoustic sensors, and thermal and gas sensors inside and around the building. The I-sensorium infrastructure has already been used in classes to implement a variety of pervasive applications/functionalities (e.g., using a mixture of video and RFID technologies to implement social policies of shared common facilities within a research building, such as reminding people to switch off the coffee machine; conduct social experiments to study recycling behavior; as well as to conduct and monitor a variety of emergency drills such as building and region evacuations). In addition, the PIs have successfully developed and deployed multiple projects using the infrastructure ranging from fire-situation awareness dashboards to learning energy behaviors in campus buildings. These prior works provide a rich application context to explore augmenting the IoT technologies with privacy preserving mechanisms. To illustrate the nature of privacy challenges and give a glimpse of the privacy technologies that could be tested and evaluated in the test bed, we discuss some initial scenarios of system usage.

Smart Building Scenario: In our instrumented space, a range of building sensors/actuators, cameras, and environmental sensors record data about people entering and exiting buildings or rooms within the buildings. Contextual information derived from these devices can be used to analyze building usage and/or occupancy levels, which can then be used to analyze/understand energy needs of a building and/or dynamic HVAC control (e.g., to reduce/shut of air conditioning in areas of buildings that are currently unoccupied). Sensor monitoring, however, may evoke

privacy concerns and is thus potentially subject to policies. Policies can be stated at different levels of abstraction corresponding to different layers of data abstraction manifested by the system. For example, an overarching policy could state:

“Do not disclose information to anyone that allows them to conclude the location of person X at a particular time with any more certainty than would be possible without the information, unless a policy of X explicitly allows it”

Very much in the spirit of recent work on differential privacy, the above policy controls disclosure by limiting what an adversary can infer using the information over and above its a-priori knowledge. A campus privacy policy could allow disclosure of classroom occupancy to any member of the campus community, but disallow access to information that would allow somebody to infer identities of people in a classroom. Surveillance and security applications require access to information such as people's location that is normally protected under the user's privacy policies. An emergency situation might warrant that information be released to appropriate authorities if the risk of not divulging such information could result in major damage to people or property. Potential information that could be derived from the gathered data and used in an emergency scenario include the following: an evacuation tally count in the case of crisis situations that result in building evacuation; location and mobility state of individuals in a region potentially affected by the hazard; and spatial messages to warn and inform individuals entering event region about potential hazards and risks.

REFERENCES

- [1] Darpa, “Brandeis.” Available: <http://www.darpa.mil/program/brandeis>. [Accessed: 17-Nov-2015].
- [2] FTC, “Internet of Things: Privacy & Security in a Connected World,” Federal Trade Commission, Jan. 2015.
- [3] D. Massaguer, S. Mehrotra, R. Vaisenberg, and N. Venkatasubramanian, “SATware: A Semantic Approach for Building Sentient Spaces,” in Distributed Video Sensor Networks, B. Bhanu, C. V. Ravishankar, A. K. Roy-Chowdhury, H. Aghajan, and D. Terzopoulos, Eds. Springer London, 2011, pp. 389–402.
- [4] “Enterprise Buildings Integrator | Honeywell Building Solutions.” <http://www.ebi.honeywell.com/en-US/Pages/homepage.aspx>. [Accessed: 17-Nov-2015].
- [5] M. Saini, P. K. Atrey, S. Mehrotra, and M. Kankanalli, “W3-privacy: understanding what, when, and where inference channels in multi-camera surveillance video,” *Multimed. Tools Appl.*, vol. 68, no. 1, pp. 135–158, Aug. 2012.