Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

SIMULATION TRUST AND THE INTERNET OF THINGS

Margaret L. Loper

Georgia Tech Research Institute
Information & Communications Laboratory
Georgia Institute of Technology
Atlanta, GA 30308 USA

ABSTRACT

The urban environment is becoming increasingly more connected and complex. In the coming decades, we will be surrounded by billions of sensors, devices, and machines, the Internet of Things (IoT). As the world becomes more connected, we will become dependent on machines and simulation to make decisions on our behalf. When simulation systems use data from sensors, devices and machines (i.e., things) to make decisions, they need to learn how to trust that data, as well as the things they are interacting with. As embedded simulation becomes more commonplace in IoT and smart city applications, it is essential that decision makers are able to trust the simulation systems making decisions on their behalf. This paper looks at trust from an IoT perspective, describing a set of research projects conducted that span multiple dimensions of trust, and discusses whether these concepts of trust apply to simulation.

1 INTRODUCTION

In the coming decade, we will be surrounded by billions of connected sensors, devices, and machines. This will lead to a pervasive presence of things (e.g., RFID tags, sensors, actuators, cell phones, vehicles), which have the ability to communicate and cooperate to achieve common goals. These things will be uniquely identifiable and addressable, and many will be smart and can capture, store, process, and communicate data about themselves, their physical environment, and their human owners. Since there is not an "internet" exclusively dedicated to "things", the expression Internet of Things (IoT) is best understood as a metaphor that encapsulates the immersion of almost anything and everything into the communications space (CDAIT 2018). As the European Research Cluster on the Internet of Things (IERC) puts it, IoT is "A dynamic global network infrastructure with self-configuring capabilities based on standards and interoperable communication protocols where physical and virtual things have identities, physical attributes and virtual personalities; use intelligent interfaces; and are seamlessly integrated into the information network" (IERC 2014).

A unique characteristics of the IoT is the presence of different modes of communication, including interaction between people (Human to Human or H2H), people and things (Human to Machine or H2M and M2H), and things (Machine to Machine or M2M). H2H communications are carried out in multiple forms and continue to innovate with social media and crowd sourcing. H2M or M2H communications assume human intervention and control. In contrast, the M2M communications have no explicit human intervention or very limited intervention.

With the growing economic and environmental problems in urban areas, the benefit of IoT technologies in a city are vast. A smart electrical grid will make cities more efficient by optimizing how energy is used and distributed. Device data will help inform and protect city residents by improving city service monitoring capabilities. Consumers will have better insights on the consumption of personal resources (energy, water, and gas) and granular neighborhood data. City infrastructures and services will change with new interconnected systems for monitoring, control, and automation (Loper 2015). Cities and urban areas that

benefit from the IoT are commonly referred to as Smart Sustainable Cities or, in short, as Smart Cities (SC): "A smart sustainable city is an innovative city that uses information and communication technologies and other means to improve quality of life, efficiency of urban operation and services, and competitiveness, while ensuring that it meets the needs of present and future generations with respect to economic, social and environmental aspects" (ITU 2014).

Application of IoT in cities include deploying *things* in all aspects of our everyday lives, including transportation systems, public safety, power grid, water supply networks, waste management, homes, buildings, health, and agriculture, as shown in Figure 1.

Figure 1: IoT and smart cities (Dhamu and Company 2019).

A goal of a smart city is to use IoT devices to collect data, and use these data to manage assets and resources efficiently, as well as to interact with the citizens that work and live in these environments. Cities may be the first to benefit from the IoT, but being surrounded by billions of sensors, devices, and machines has profound implications for security, trust, and privacy. The more technology a city uses, the more vulnerable it is, so the smartest cities face the highest risks. Therefore, in addition to using IoT technologies for monitoring and managing cities, there is also a need to understand how to protect and defend smart cities.

This paper describes a portfolio of research conducted on trust as it applies to IoT and SC. It starts by defining trust and a framework for capturing the many dimensions of trust. That is followed by brief descriptions of some of the research projects conducted on trust in IoT and SC. These projects are organized in terms of the trust framework we adopted. The research described spans a range of trust topics, but little work has been done on the relevance of trust to simulation. The application of simulation that use and interact with real world systems is growing, which means understanding their trustworthiness is of growing importance. The paper concludes with a few ideas on the way forward.

2 TRUST

2.1 Definitions

Several trends are emerging with IoT. First, with the grand vision of billions and trillions of things (e.g., cellphones, physical devices, vehicles, and other things embedded with electronics, software, sensors, actuators), things will soon outnumber people. One prediction forecasts that by 2021, the number of

connected devices will outnumber connected people by six to one (ITU 2012). It will be impossible for humans to monitor and control all these things; therefore, some decision-making will be delegated to things in the system. In other words, some of these things will make decisions on our behalf. Second, while interconnection is self-evident to IoT, the intelligence of things is what makes the IoT paradigm "game-changing" (CDAIT 2018). There is an increasing desire to use things in lieu of humans in dangerous or routine situations, and also to make things more intelligent such that they can deliver personalized and autonomic services. Both of these trends raise questions about the trustworthiness of this emerging technology.

The connection between people and things is complex, and creates a set of trust concerns. Trust should be considered at two levels: (1) whether a thing trusts the data it receives or trusts the other things it interacts with (M2M) and (2) whether a human trusts the things, services, data, or IoT offerings that it uses (H2M or M2H). This leads to the idea that trust is multi-dimensional. Ahn et al. (2007) described the concept of multi-dimensional trust by different agent characteristics, such as quality, reliability, and availability. For Matei et al. (2009), trust refers to the trustworthiness of a sensor, whether it has been compromised, the quality of data from the sensor, and the network connection. Grandison and Sloman (2000) define trust as the belief in the competence of an entity to act dependably, securely, and reliably within a specified context. To address behavior uncertainty in agent communities, Pinyol and Sabater-Mir (2013) define three levels of trust based on human society: security, institutional, and social. Leisterm and Schultz (2012) identify technical, computational, and behavioral trust, but focus primarily on a behavioral trust indicator. Lastly is the idea that trust is a level of confidence: the probability that the intended behavior and the actual behavior are equivalent given a fixed context, fixed environment, and fixed point in time (Voas et. al 2018).

For our work, we adopted the definition of trust that NIST uses in their report on trustworthiness of cyber physical systems. Trust is defined as "… the demonstrable likelihood that the system performs according to designed behavior under any set of conditions as evidenced by characteristics including, … security, privacy, reliability, safety and resilience" (NIST 2017, p.15).

2.2 Types of Trust in IoT Systems

Work on trust management is often divided into two areas: security-oriented and non-security- oriented. The descriptions below are summarized from (Terzis 2009).

Security-oriented trust adopts a restricted view, where trustworthiness is equated to the degree to which an entity or object is considered secure. This traditional view sees trustworthiness as an absolute property that an entity either has or doesn't have. This is often accomplished by determining the credentials an entity possesses, and iteratively negotiating how to disclose certified digital credentials that verify properties of trust. This view of trust is also related to trusted computing, which is the expectation that a secure operating environment can be created by enforcing certain hardware and software behaviors with a unique encryption key inaccessible to the rest of the system. In software engineering, this view of trust is determined through formal verification. Managing trust in this context includes specifying and interpreting security policies, credentials, and relationships.

Non-security-oriented trust adopts a wider view similar to the social sciences. This includes a view of trust as a mechanism for achieving, maintaining, and reasoning about the quality of service and interactions. In this view, trust is a measurable property that different entities have in various degrees. Trust is determined on the basis of evidence (personal experiences, observations, recommendations, and overall reputation) and is situational, meaning an entity's trustworthiness differs depending on the context of the interaction. A goal of trust management is managing the risks of interactions between entities. This is also the basis of trust management in multiagent systems, which includes the notion of malicious and selfish behavior. Since non-security-oriented trust is similar to the human notion of trust, work related to computer-mediated trust between users, building human trust in computer systems, and human-computer interaction has led to sophisticated models of trust and reputation research.

To tie this together in a system model for IoT, we adopt a layered trust framework defined by Yan et al. (2014). These layers work together to create an environment in which things and humans can interact

and make trustworthy decisions. The layers in the framework include (i) physical perception, which perceives physical environments and human social life; (ii) a network layer that transforms and processes perceived environment data; and (iii) an application layer that offers context-aware intelligent services in a pervasive manner. The fourth layer represents the cyber-physical social relationships that connect layers. Figure 2 depicts these layers, with trust objectives. A trustworthy IoT system relies on the cooperation among layers. "Ensuring the trustworthiness of one IoT layer (e.g., network layer) does not imply that the trust of the whole system can be achieved" (Yan et al. 2014).

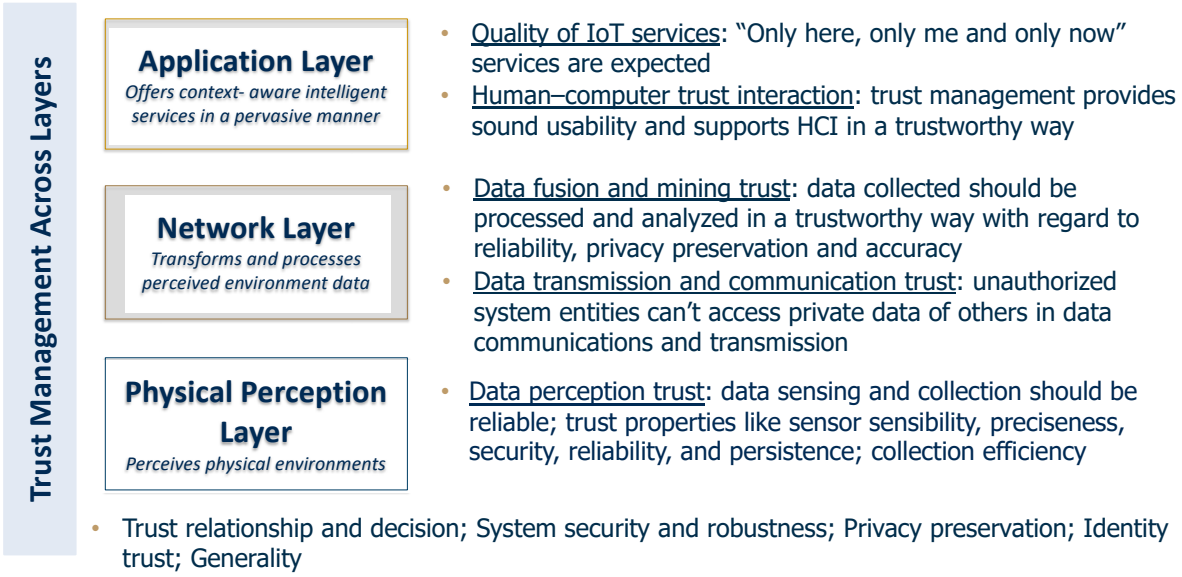

Figure 2: IoT trust framework.

2.3 Trust Architecture

The next step we took is to translate the layers of the trust framework into an architecture, on which a research strategy was developed. The trust architecture and its system components are shown in Figure 3a. The lowest component called sensor is analogous to the physical perception layer, where sensors, devices, and machines are individually serving as a source for data. The data being generated by these sensors are assimilated and elevated through context into information by an aggregator.

An aggregator is an intelligent machine that collects data from sensors and uses that data to create knowledge for decision-making. In order for an aggregator to determine if the data communicated to it are worthy of being used, the notion of trust becomes an issue. When there are a number of different sensors, each presenting data, and these data have conflicts from sensor to sensor, aggregators must select which among them to trust, how much to trust, as well as some criteria for establishing that trust. Reasons for competing sources of data to be in conflict with each other, at the Aggregator or Command and Control (C2) components, could be many: malfunction, bad actor, tampering, environmental conditions, context conditions, and so on. Finally, the C2 component is responsible for looking across aggregators to synthesize data, as well as provide an interface to humans interacting with the system.

3 TRUST RESEARCH STRATEGY

Our research in IoT started in 2013 with an ideation event that engaged a large number of researchers to discuss technologies, technical challenges, and application areas. This was followed by a number of internally-.funded research projects, shown in Figure 3b, which spanned the trust architecture. Our initial research was on IoT M2M Trust, which looked at how aggregators trust data they receive from sensors, as

well as how they trust other aggregators. That research was followed by trust and security of sensor arrays, and a trust negotiation language that could be used by aggregators. Our research broadened into smart cities at this point, where we looked at computing trust at the edge in Fog Computing (defined later), in-situ privacy algorithms, intelligent and scalable orchestration of IoT objects, and a situational awareness system for the Georgia Tech police department called the Common Operating Picture. In recent years, we have looked at the question of risk factors to determine if smart cities are trustworthy, as well as modeling the inherent danger associated with smart city technology being manipulated.

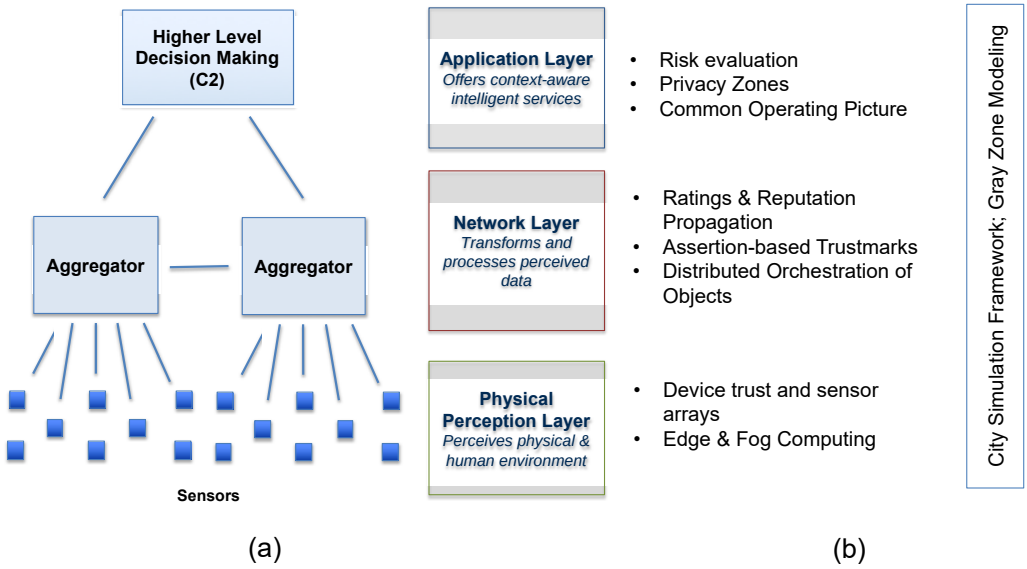

Figure 3: (a) Trust architecture and (b) research strategy.

In the sections that follow, many of these projects are described to give a better understanding of our trust research strategy. Note that these projects span the multiple dimensions of trust defined in Section 2.1, as well as security- and non-security-oriented trust approaches.

3.1 Physical and Perception Layer

3.1.1 Device Trust and Sensor Arrays

With a growing demand for data collection, there is also a demand for privacy and safety. This project developed a sensor testbed and algorithms to collect scientific research data, while preserving privacy, in public environments. The technical approach interfaced a reconfigurable System-on-Chip with multiple sensing peripherals and wireless network feedback. The sensor arrays, Community Array Nodes, contained ten different sensors that gave insight into urban environment monitoring and smart city behavior. Hardware- and software-based security and trust mechanisms allowed for operating in remote environments, providing protection to both deployed hardware and the back-end infrastructure.

There were three successful hardware deployments across the city of Atlanta's North Avenue smart city test bed, each with the capability of collecting sensor data, processing it in-situ, and sending it to a central server over a secure connection. With only an outgoing connection from the node, there were minimized avenues for security threats. Hardware security was accomplished through tamper-resistant technology, i.e., resistance to intentional malfunction or sabotage by normal users or others with access to the technology. The reconfigurable sensing platform demonstrated the ability to effectively collect relevant

information, while alleviating public concerns. Precedent-aware classification (Danner 2016), an efficient technique that uses information from past detections to reduce the computation required for detection over time, allowed computationally intensive computer vision algorithms to run in real time. With a modular design of the nodes, they were quickly integrated with relevant sensors from researchers. This approach minimized the difficulty in deploying a distributed sensor network, while addressing privacy and security concerns.

3.1.2 Fog Computing

Due to its widespread availability, high degree of scalability, and inexpensive cost, cloud computing is ubiquitous in industry. Despite its extreme versatility, cloud computing cannot address all possible computing needs for modern day applications. For example, if the network connection to the Internet is severed, the uplink too slow, or the distance between the user and the server too great, the user's cloud experience will be significantly degraded. To that end, Cisco proposed a distributed computing framework called *Fog* Computing (Bonomi et al. 2012). Fog Computing ameliorates these issues by creating a continuum of compute, storage, and networking resources from the cloud to the end device. Fog Computing primarily targets applications which have high bandwidth requirements, low latency restrictions, or that operate in environments where internet and cloud connectivity are restricted, denied, or intermittent. In this work, we explored Fog Computing's ability to operate under degraded networking conditions, which addresses the reliability aspects of trust.

We developed a real-world Fog Computing application that processed multiple, collocated video streams to extract statistical information about how pedestrians use walking paths. To develop the application, we used Georgia Tech's MobileFog platform (Hong et al. 2013; Saurez-Apuy et al. 2016) in conjunction with data from the PETS2009 (Ferryman et al. 2010) and Duke Multi-Target Multi-Camera (Ristani et al. 2016) data sets. Using the MobileFog framework, we developed a pedestrian statistics application that located human figures in frames, tracked their movement between frames, and generated ground position estimates with decimeter-level accuracy. These data were transmitted up through the Fog and twice aggregated before reaching the root node where they were logged to disk. Our experiments tested the limitations of Fog Computing in three bandwidth configurations: benign, hostile, and denied. Our results represent the first systematic exploration of a real-world Fog Computing application's response to degraded networking conditions.

3.2 Network Layer

3.2.1 Trust Negotiation Language

A trust framework is any structure that builds trust among autonomous actors for the purpose of sharing and reusing identities. The goal of the Trustmark framework (GTRI 2013) is to facilitate federated identity and attribute management (i.e., the reuse of digital identities and associated attributes) in enterprise systems. Identity reuse requires trust between entities that assert attributes and entities that rely on such assertions. The rules and requirements for establishing such trust comprise an identity trust framework. The requirements of a trust framework may be explicitly or implicitly stated, and may encompass many dimensions such as identity assurance, privacy, security, technical interoperability, business-level identity requirements, legal rights, responsibilities, liabilities, and indemnification.

This project revolved around the use of Trustmarks as a secure and robust framework for exchanging trusted, third-party-attested attributes in support of autonomous peer-to-peer trust decisions. A Trustmark is a machine-readable, cryptographically signed digital artifact, issued by a Trustmark provider to a Trustmark recipient, and relied upon by one or more Trustmark- relying parties. Such a process is valuable in IoT, as devices come from manufacturers that are themselves Trustmark recipient organizations. These IoT-enabled devices can then present preloaded Trustmarks in order to establish trust. For example, if the

device manufacturer Samsung created smart phones with a particular Android capability, it would make sense for Android to grant a Trustmark to Samsung which would then be preloaded and presentable on all Samsung-made devices, to prove that they adhere to a particular set of requirements laid forth by Google on the Android platform.

To securely exchange attribute information in support of IoT trust and trust negotiation, we developed a set of extensions to a pre-existing Trustmark framework specification, to include parameter definitions and values within Trustmarks. These parameters contain data-typing information that allows for conveying most modern attribute information, as well as human-readable names and descriptive text for helping Trustmark assessors fill in the appropriate values before issuing Trustmarks. Within Trustmarks, these parameters provide the necessary third-party-attested values that are required for rich trust decisions.

3.2.2 M2M Trust

Machine to machine communications are at the center stage of the IoT. Connecting the physical world with the digital world not only creates new opportunities for innovation and discovery, but also opens doors for misuse and abuse. This work argues that reputation-based trust can be an effective countermeasure for securing M2M communications. We established M2M trust by taking into account both transaction/interaction service behaviors and feedback-rating behaviors in the presence of bogus transactions and dishonest feedback. Our trust model, called M2MTrust (Liu et al. 2016), introduces two novel trust metrics: pairwise-similarity-based feedback credibility and threshold-controlled trust propagation. We compute the direct trust from machine A to machine B by utilizing their pairwise rating similarity as the weight to the normalized aggregate of ratings that A has given to B.

We examined the strength and weakness of several popular trust models developed in the context of decentralized network computing systems, such as EigenTrust (Kamvar et al. 2003), PeerTrust (Xiong and Liu 2004), BetaTrust (Jasong and Ismail 2002), and ServiceTrust (Zhiyuan et al. 2015). We evaluated these trust models in terms of three sets of measurements:

1. The support of only direct trust evaluation vs. the support of both direct trust and indirect trust;
2. The time complexity for efficiency in trust computation and trust deployment in M2M communication;
3. The attack resilience against four common threat models initially introduced in (Kamvar et al. 2003): malicious individuals, malicious collective, camouflaged collective, and malicious spies.

We conduct extensive experiments using simulation and real datasets. Our scenario was self-driving cars on road networks. Specifically, can self-driving cars trust one another to provide a safe driving experience, and can M2MTrust help alleviate traffic jams, whether accidental or malicious? Our direct trust computation model effectively constrained malicious nodes to gain direct trust from dishonest feedback ratings by leveraging feedback credibility. Furthermore, our threshold-controlled trust propagation mechanism successfully blocked the trust propagation from good nodes to malicious nodes. Experimental results showed that M2MTrust significantly outperformed other trust metrics in terms of both attack resilience and performance in the presence of dishonest feedback, and sparse feedback ratings against four representative attack models.

3.2.3 Distributed Orchestration in Large-Scale IoT Systems

With the growing popularity of smart things and the pervasiveness of wireless communications, the Internet has evolved from the Internet of hosts, to the Internet of people, to the IoT. Intelligent and scalable orchestration of large-scale IoT objects using a multitier architecture is critical to embrace the vision of IoT. In Yigitoglu et al 2017), we present our vision and our initial development of a distributed orchestration framework, called ISYMPHONY, with the ultimate goal of scaling real-time and on-demand IoT service

provisioning in large-scale IoT systems, while guaranteeing quality of service with respect to performance, availability, and security.

This research made two original contributions. First, we presented a distributed orchestration architecture that enabled edge devices and edge clients running on top of edge devices to contribute to the IoT computing tasks, based on their computing and communication capacities. A main idea behind the distributed IoT orchestration architecture was the intelligent partition of a real-time IoT computing task into an optimal coordination of server-side processing and IoT-object-side processing. Second, a set of optimization techniques limit the number of computations handled by the edge clients and enhance the overall performance and resource utilization of the ISYMPHONY system. Our initial experimental results show that ISYMPHONY can lead to significant savings in terms of server load and high accuracy by leveraging and coordinating edge client processing capabilities, compared to the solutions relying solely on server level processing for real time IoT services provisioning.

3.3 Application Layer

3.3.1 Trusting Smart Cities

The benefits of making cities smart must be considered against the potential harm that could come from being massively interconnected. To understand how trust applies to smart cities, we developed a set of risk factors that capture a range of issues that cities should considere when deploying smart city technologies (Loper 2018). Risk emerges when the value at stake in a transaction is high, or when this transaction has a critical role in the security or the safety of a system. “In most trust systems considering risk, the user must explicitly handle the relationship between risk and trust by acknowledging that the two notions are in an inverse relationship, i.e. low value transactions are associated to high risk and low trust levels and vice versa” (Patrick 2002). Or put more simply – the more risk associated with a transaction, the less we trust it.

To understand how this inverse relationship applies to smart cities, we defined three key risk factors (Loper 2015): non-technical, technical, and complexity. Non-technical risk includes aspects of a smart city where humans are involved, such as management, training and education, governance, and security practices. Technical risk factors focus on the technology aspects of a smart city, including both hardware and software systems. This also includes the concept of cyber-physical systems, which are systems of collaborating computational elements controlling physical entities. The last risk factor is complexity. A smart city is not a discrete thing; it is the complex multi-dimensional interconnection of diverse systems (human and technology) that deliver services and promote optimum performance to its users. There is risk in the complexity of these systems, especially as the scale becomes very large. Building on these risk factors, a threat analysis matrix for capturing how well smart cities address these risks was proposed.

3.3.2 Privacy Zones

While location-based services and applications are increasing in popularity, there are growing concerns over users’ location privacy. Although there exist general-purpose mobile permission systems and cloaking techniques, they suffer from several problems when applied to continuous location and GPS access. Namely, they are often rigid, coarse-grained, not sufficiently personalizable, and unaware of road network semantics. For example, in most existing systems, permission decisions are static and follow a one-size-fits-all principle. Once a user decides to allow or deny GPS access to an app, the setting is applied on all future location requests unless the user manually changes the app’s setting. In Yigitoglu et al. (2018) we proposed PrivacyZone, a novel system for constructing personalized fine-grained privacy regions and protecting users’ privacy within these regions. PrivacyZone allows users to seamlessly enter their privacy specifications under spatial, temporal, and semantic customization. For example, a user may want to allow location access when she is in the park, but deny access when she is in the hospital. Existing permission mechanisms are not sufficiently fine-grained to support such personalized and variable privacy preferences.

Novel challenges arise from enforcing privacy zones for a large volume and variety of users with frequent location updates. We show that I privacy zone processing techniques are inefficient and cause excessive energy consumption. As a result, we developed advanced processing techniques based on the concept of safe hibernation – a time period or a geographic region, within which the client is guaranteed to not enter a privacy zone. We empirically evaluated our techniques to demonstrate their trade-offs with respect to hibernation time, computation effort, and network bandwidth usage. Our results show that PrivacyZone is efficient, scalable, and flexible, while preserving the users' location privacy.

4 SIMULATION TRUST

The paper thus far has done the following: converged on a definition of trust, defined a framework that captures its multiple dimensions, and investigated a number of research projects that span these layers. The next step is to understand whether these concepts of trust apply to simulation.

4.1 LVC and IoT

As mentioned earlier, IoT is characterized by a wide variety of tags, sensors, actuators, analytics, and embedded systems that are uniquely identifiable and addressable, and cooperate over networks. Also discussed, IoT has multiple modes of interaction, which include people and things (e.g., H2H, H2M, M2H, and M2M). Characterizing IoT interactions between people and things resembles the framework developed for how people and simulation models interact, known as Live, Virtual, and Constructive (LVC) (MSCO 2011). The similarities between these two areas are quite interesting, and lead us to the observation that IoT is another type of LVC system. The LVC taxonomy, shown in Figure 4, is defined as:

- Live simulation refers to M&S involving real people operating real systems (e.g., a pilot flying a jet).
- Virtual simulation is one that involves real people operating simulated systems (e.g., a pilot flying a simulated jet).
- Constructive simulation applications are those that involve simulated (or no) people operating simulated systems (e.g., a simulated pilot flying a simulated jet).

There is no name for simulated people operating real equipment. In the late 1980's when the LVC taxonomy was created, there were no examples of this type of interaction. However, technology has advanced to the point where simulated humans are operating real systems. For example, driverless cars have proved that the interaction between the real and simulated worlds is possible. Even though that quadrant of the matrix has not been officially named, it bears resemblance to artificial intelligence and autonomy.

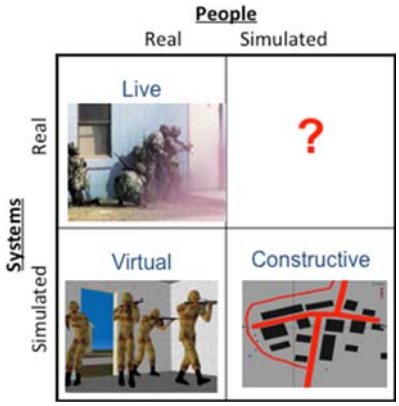

Figure 4: Categorizing simulation models by the way humans interact with them (IITSEC 2018).

First presented in (Loper 2017), we can use the LVC taxonomy to describe IoT, as shown in Figure 5.

- Live refers to real people operating real IoT systems (e.g., a smart phone).
- Virtual refers to real people operating simulated IoT systems (e.g., social media).
- Constructive refers to simulated (or no) people operating simulated IoT systems (e.g., analytics).
- Autonomy refers to simulated (or no) people operating real IoT systems (e.g., driverless vehicles).

	Simulation	IoT
Live	Real people operating real systems	Sensors, Devices, Smart Phones, Security Cameras
Virtual	Real people operating simulated systems	Mobile Apps, Social Media, Driving Directions
Constructive	Simulated (or no) people operating simulated systems	Embedded Simulation, Machine Learning, Analytics
Autonomy	Simulated (or no) people operating real systems	Driverless Vehicles, Robots, Home/Building Automation, Embedded Systems

Figure 5: Using LVC taxonomy for IoT.

The map between IoT and LVC in Figure 5 helps us think about how trust applies to simulation. Using the projects discussed in Section 3 as a roadmap, we can posit whether trust research applies to live, virtual, or constructive simulation. Let's explore several examples: First, the device trust for sensor arrays and privacy zone work could be applied to live simulation. Context is often necessary for representing things in simulation models (e.g., location), however knowing the data are different than knowing the person(s)/vehicle(s) that generated the data. This highlights the need to protect the privacy of the data collected from sensors. A second area applies to one of the fundamental elements of LVC simulation: communication mechanisms (exchanging data). The work on M2MTrust and Trustmarks could easily apply to the messages exchanged by simulation system, as well as determining which simulation model to trust in a message exchange. Third, the reliability of Fog Computing has direct application to live simulation, and constructive in regards to simulation in the cloud or simulation as a service. When simulation computations are pushed closer to the edge to improve response time, the trustworthiness of these computing platforms is critical. Lastly, the risk framework created for looking at smart cities as a whole is directly relevant to looking at the development and execution of LVC+A federations, when one or more of the simulation components are driven by live sensor data. From this quick analysis, it appears that the portfolio of trust research we have explored has relevance to the continuum of LVC simulation.

4.2 Internet of Simulation Things

IoT enables distributed control and computational architectures: one can trust (or distrust) abstract concepts, abstract entities, or physical things; including persons, organizations, information, systems, etc. Since simulation models can be used to control or give commands to sensors and actuators, or provide faster-than-real-time prediction to systems, we need to enhance trust relationships when simulation is part of the IoT system. Expanding IoT's modes of interaction, we have:

- Machine to Machine (M2M) ➔ Simulation to Simulation (S2S) or Machine to Simulation (M2S)
- Human to Machine (H2M) ➔ Human to Simulation (H2S)
- Human to Human (H2H) ➔ Human to Human (H2H)

An example of where M2S and S2M are already happening is data-driven online simulation. The Dynamic Data-Driven Application Systems (DDDAS) concept is a unique paradigm for exploiting

maturing computational and sensor networking technologies to compensate for model deficiencies and unforeseen system evolution and stimulus conditions, mitigate the effect of design imperfections on long-term as well as short-term system safety, and enable informed decision for maintenance planning and crisis management (Farhat et al. 2006). This paradigm utilizes online data to drive simulation computations, and the results are then used to optimize the system or adapt the measurement process. For example, live sensor data and analytics can be used to construct or infer the current state of a system and faster-than-real-time simulation can then be used to project the system's future state. Also, simulation can be used to control an operational system, e.g., data from a real system are fed directly into the simulation model which analyzes alternate options and produces recommended courses of action. With the availability of data from IoT and smart city instrumentation, paradigms such as DDDAS can be expected to grow in importance.

As discussed by Carothers et al. (2017), when simulation uses data from *things* in the network to make decisions, users need to learn how to trust these data as well as the things (sensors) they are interacting with. Currently securing sensors and devices is accomplished through information security technologies, including cryptography, digital signatures, and electronic certificates. This approach establishes and evaluates a trust chain between devices, but it does not tell us anything about the quality of the information being exchanged over time. Data from sensors or aggregators may be in conflict with each other due to malfunction, bad actors, tampering, environmental conditions, context conditions, and so on. Thus, whether or not the simulation should trust these data must be established by an agent that is capable of a trust evaluation prior to them being deemed useful as information. Further, if simulation has a role in controlling or giving commands to some sensor or actuator in the IoT system, then the data the simulation uses from external sources in which to make those decisions must be trustworthy, such that they are not purposely misled into issuing malicious commands.

To illustrate the emerging importance of trust for simulation, let's first look at where simulation might be used in smart cities. An ontology that represents a city as a system of systems is the Anatomy of a City, developed by the City Protocol Society (City Protocol Society 2015), shown in Figure 6. This document defines a common language describing the city ecosystem as three key system elements: Structure – a set of physical structures; Society – the living entities that make up a city's society; and Interactions – the flow of interactions between structure and society.

The top layer of the ontology is Structure, which refers to physical constructions in a city, i.e., the building, streets, subways, and other three-dimensional macroscale networks. This layer, which is particularly useful in the trust discussion, includes the following sub-layers:

- *Environment* is the physical and geographic setting of the city, including the natural environment ("nature"). It is formed by the three basic elements: air, earth, and water, which interact dynamically in a seasonally-dependable way.
- *Infrastructure* comprises the connective structures that enable people to get resources, especially from the environment, and bring them to the city, or that enable flows or cycles within the city itself. These infrastructures include those that support communications, the water and energy cycles, the matter cycle that supports the movement of goods and food as well as the resultant waste, the mobility networks, and nature or green infrastructure.
- *Built Domain* is organized according to the approximate number of people that it can accommodate on a physical basis. Within the Built Domain, an object corresponds to a single person, and a dwelling, building, block, neighborhood, district, city and metropolis or region each increase the scale by an order of magnitude. Private and public spaces are contained within each level of scale.

Work is already underway to use simulation to monitor, control, and predict aspects of cities. Related to the built domain sub-layer, Farhat et al. (2006) are using a DDDAS to monitor the health of large-scale structural systems. Their work is focused on composite materials of aircraft, but we can envision it being applied to city structures like stadiums, bridges, or dams. The overall goal for their work is to enable and promote active health monitoring, failure prediction, aging assessment, informed crisis management, and

decision support for complex and degrading structural engineering systems. Related to the infrastructure sub-layer, simulation and optimization can be used to monitor a city's water supply. Mahinthakumar et al. (2006) recognize that urban water distribution systems are vulnerable to accidental and intentional contamination incidents that could result in adverse human health and safety impacts. When a contamination event is detected, e.g., data from a water quality surveillance sensor network and reports from consumers, they use a DDDAS approach to answer critical questions like what response action (e.g., shut down portions of the network, implement hydraulic control strategies, introduce decontaminants) should be taken to minimize the impact of the contamination event. Real-time answers to complex questions can be addressed through dynamic integration of computational components (including models and simulation) and real-time sensor data. The last example is also related to infrastructure, focused on transportation. In Saroj et al. (2018), a real-time data-driven transportation simulation model was used to evaluate and visualize network performance, and provide dynamic operational feedback. The study used a hybrid traffic simulation model to represent seventeen consecutive intersections on a traffic corridor partially equipped with smart devices. The architecture would enable control of the signals and the vehicle volumes using real-time data from in-field detectors.

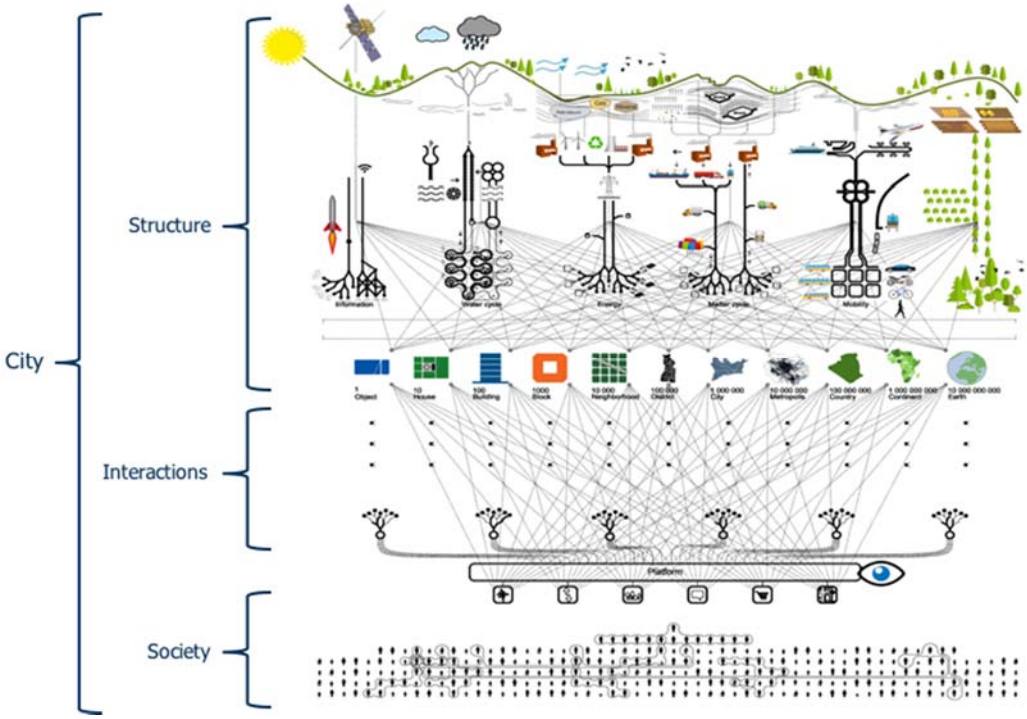

Figure 6: Anatomy of a city (City Protocol Society 2015).

As more data-driven simulation is used in smart cities, a concern is that the sensors that feed data into the simulation systems can be hacked and fed fake data. This could be used for all manner of mischief, like causing signal failures that shut down subways or allowing contaminants into the water supply. For example, what if the data driving transportation simulation systems made traffic signals stay red or green, tweak electronic speed limit signs, or messed with ramp meters to send cars onto the freeway all at once? What would commuting look like if erroneous sensor data sent to simulation changes the routes of public transportation or changes subway schedules? How would cities respond to an inadequate supply of electricity or water, or worse yet, not be notified that drinking water was contaminated? What if waste collection was interrupted during the summertime, and garbage piled up in the streets because the data from smart trash cans that feed a simulation to optimize trash routes was misrepresented? Many systems in cities

are interconnected, so erroneous data driving one simulation could cause a cascade effect, impacting other systems in the city.

Many of these issues get at data integrity, and how to detect misbehavior in the sensor system. Too many false positives may remove valuable sensor resources from the network, while too many false negatives may pollute the data generated and veer the simulation off track (Farhat et al. 2006): Research that looks at the sensor networks that drive simulation models, and how to discover and correct node misbehavior is critical for simulation trust.

5 CONCLUSIONS

As the number of sensors and simulation applications connected to the network grows, we will see different patterns of communication and trust emerge. Data from the sensors and higher-level aggregators will be fed into models and simulation models that are making predictions and decisions that will impact our lives. Creating, understanding, and managing large-scale distributed simulation systems interacting with each other to manage operational systems present major challenges. As pervasive simulation becomes more commonplace in IoT and smart city applications, it is essential that it is secure or at least tolerant of cyber threats. Privacy and trust issues must also be adequately addressed to realize widespread adoption.

This paper has covered the topic of trust in IoT and smart cities, and posed an argument for why this work is directly relevant to simulation. Future research should focus on fundamental principles concerning how trust is established, maintained, and used in simulation, and a theory behind their operations. Simulation validation is part of the solution to this problem, and work such as model trust through curation is also starting to look at parts of the space. However, simulation trust is not an area that simulation research has traditionally focused. A valid question is whether this work should come from the simulation community, or whether it belongs to the cyber security community instead.

The definition of trust has many dimensions, which means that there is a rich landscape of problems to address. To construct a set of research issues to consider for simulation trust, we can look to a recently published NIST report which identifies 17 technical concerns that negatively affect the ability to trust IoT products and services (Voas et al. 2018):

1. Scalability
2. Heterogeneity
3. Ownership and Control
4. Composability, Interoperability, Integration, and Compatibility
5. “Ilities” (availability, compatibility,…)
6. Synchronization
7. Lack of Measurement
8. Predictability
9. Testing and Assurance
10. Certification
11. Security
12. Reliability
13. Data Integrity
14. Excessive Data
15. Speed and Performance
16. Usability
17. Visibility and Discovery

Some of these concerns – testing and assurance, certification, heterogeneity, interoperability, composability – are areas where the simulation and LVC community has spent considerable time developing solutions. In other areas – reliability, data integrity – we have spent less time. All of these factors, plus insurability and risk measurement, represent new areas of research that we should pursue to ensure simulation trust in untrusted environments.

ACKNOWLEDGMENTS

I would like to thank the Georgia Tech Research Institute and the Georgia Tech Institute for Information Security and Privacy for supporting this research. This work could not have been accomplished without my collaborators in GTRI and in the Georgia Tech College of Computing. I appreciate their collaboration and contributions over the last six years.

REFERENCES

Ahn, J., D. DeAngelis, and S. Barber. 2007. "Attitude Driven Team Formation Using Multi-Dimensional Trust". In *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology* (IAT '07), 2nd–5th November, Fremont, CA, 229 –235.

Bonomi, F., R. Milito, J. Zhu, and S. Addepalli. 2012. "Fog Computing and its Role in the Internet of Things". In *Proceedings of the First MCC Workshop on Mobile Cloud Computing*, 13th–17th August. New York, NY, 13–16.

Carothers, C., A. Ferscha, R. Fujimoto, D. Jefferson, M. Loper, M. Marathe, S. Taylor, and H. Vakilzadian. 2017. "Computational Challenges in Modeling and Simulation". In *Research Challenges in Modeling & Simulation for Engineering Complex Systems*, edited by R. Fujimoto, C. Bock, W. Chen, E. Page, and J. H. Panchal, 45–74. Heidelberg: Springer Nature.

CDAIT. 2018. *Driving New Modes of IoT-facilitated Citizen/User Engagement*. Center for the Development and Application of Internet of Things Technologies, Technical Report, July 2018. https://cdait.gatech.edu/sites/default/files/georgia_tech_cdait_thought_leadership_working_group_white_paper_july_9_2018_final.pdf.

Dhamu and Company. 2019. *Smart City Jaipur & Udaipur Get Rs. 345 Crore from Government*. https://www.dhamuandcompany.net/property-jaipur/smart-city-jaipur-udaipur-get-rs-345-crore-from-government/, accessed 18th August.

City Protocol Society. 2015. *City Anatomy: A Framework to Support City Governance, Evaluation and Transformation*. City Protocol Agreement (CPA-I_001-v2), https://cpsociety.sharepoint.com/sites/cptf/CPTSC/Private%20Documents/Publications/CPA-I_001-v2_City_Anatomy.pdf, accessed 23rd August.

Danner, J. 2016. "Rapid Precedent-aware Pedestrian and Car Classification on Constrained IoT Platforms", *ESTIMedia Conference*, 13th–18th October, New York, NY.

Farhat, C., J. G. Michopoulos, F. K. Chang, L. J. Guibas, and A. J. Lew. 2006. "Towards a Dynamic Data Driven System for Structural and Material Health Monitoring" In *International Conference on Computational Science*, edited by V. N. Alexandrov, G. D. van Albada, P. M. A. Sloot, and J. Dongarra, 456-464. Berlin, Heidelberg: Springer.

Ferryman, J. and A. Shahrokni. 2009. PETS2009: Dataset and Challenge. In *2009 Twelfth IEEE International Workshop on Performance Evaluation of Tracking and Surveillance*, December 7th–9th. Snowbird, UT, paper 1–6. https://doi.org/10.1109/PETS-WINTER.2009.5399556.

GTRI. 2013. *NSTIC Trustmark Pilot*. Georgia Tech Research Institute. Available from: https://trustmark.gtri.gatech.edu.

Grandison T. and M. Sloman. 2000. "A Survey of Trust in Internet Applications". *IEEE Communications Surveys and Tutorials* 3(4): 2-16.

Hong, K., D. Lillethun, U. Ramachandran, B. Ottenwälder, and B. Koldehofe. 2013. Mobile Fog. In *Proceedings of the Second ACM SIGCOMM Workshop on Mobile Cloud Computing* (MCC'13). 12th–16th August, Hong Kong, China. https://doi.org/10.1145/2491266.2491270.

IERC. 2019. *IoT European Research Cluster Website*. http://www.internet-of-things-research.eu/about_iot.htm, accessed 17th June.

IITSEC. 2018. "Fundamentals of Modeling and Simulation". *Interservice/Industry Training, Simulation and Education Conference* (I/ITSEC). 26th November, Orlando, FL, Tutorial Number 1819.

ITU. 2012. "The State of Broadband 2012: Achieving Digital Inclusion for All". International Telecommunication Union, Broadband Commission for Digital Development, Technical Report, https://www.itu.int/pub/S-POL-BROADBAND.04, accessed 23rd August.

ITU. 2014. *Smart Sustainable Cities: An Analysis of Definitions.* International Telecommunication Union, Focus Group on Smart Sustainable Cities (FG-SSC). https://www.itu.int/en/ITU-T/focusgroups/ssc/Documents/website/web-fg-ssc-0100-r9-definitions_technical_report.docx, accessed 23rd August.

Jasong A. and R. Ismail. 2002. "The Beta Reputation System". In *Proceedings of the 15th Bled Electronic Commerce Conference*, 17th–19th June, Bled, Slovenia.

Kamvar, S. D., M. T. Schlosser, and H. Garcia-Molina. 2003. "The Eigentrust Algorithm for Reputation Management in P2P Networks". In *Proceedings of the 12th International Conference on World Wide Web*, 20th–24th May, Budapest, Hungary, 640–651.

Leisterm, W. and T. Schultz. 2012. "Ideas for a Trust Indicator in the Internet of Things". In *Proceedings of the First International Conference on Smart Systems, Devices and Technologies* (SMART 2012), 27th May – 1st June, Stuttgart, Germany, 31-34.

Liu, L.,M. Loper, Y. Ozkaya, Y. A. Yasar, and E. Yigitoglu. 2016. "Machine to Machine Trust in the IoT Era". In *The 18th International Workshop on Trust in Agent Societies* (Trust 2016), 9th–13th May, Singapore, 18–29.

Loper, M. 2015. *Trusting Smart Cities*. Technical Report, Sam Nunn Security Fellows Program, Georgia Institute of Technology, Atlanta, Georgia.

Loper, M. 2017. "Trust as a Service in LVC Simulations". Invited Panel on Research Challenges in M&S in the Era of Big Data and the Internet of Things, *Interservice/Industry Training Simulation & Education Conference* (I/ITSEC), 27th November – 1st December, Orlando, FL.

Loper, M. 2017. "Trusting Smart Cities: Risk Factors and Implications". In *Small Wars Journal*, presented at the TRADOC Mad Scientist Conference on Installations of the Future, June 19th, Atlanta, GA. http://smallwarsjournal.com/jrnl/art/trusting-smart-cities-risk-factors-and-implications.

Mahinthakumar, K., G. von Laszewski, R. Ranjithan, D. Brill, J. Uber, K. Harrison, S. Sreepathi, and E. Zechman. 2006. "An Adaptive Cyberinfrastructure for Threat Management in Urban Water Distribution Systems". In *International Conference on Computational Science*, Berlin, Heidelberg: Springer, 401-408.

Matei, I., J. Baras, and T. Jiang. 2009. "A Composite Trust Model and its Application to Collaborative Distributed Information Fusion". In *Proceedings of the 12th International Conference on Information Fusion* (FUSION 2009), 6th–9th July, Chicago, IL, 1950 –1957.

MSCO. 2011. *DoD Modeling and Simulation (M&S) Glossary. https://www.msco.mil/MSReferences/Glossary/MSGlossary.aspx*.

NIST. 2017. National Institute for Standards and Technology, *NIST Special Publication 1500-202 Framework for Cyber-Physical Systems*: Volume 2, Working Group Reports, Version 1.0, June 2017, http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-202.pdf.

Patrick, A. 2002. *Building Trustworthy Software Agents*. IEEE Internet Computing 6(6):46–53.

Pinyol I. and L. Sabater-Mir. 2013. "Computational Trust and Reputation Models for Open Multi-Agent Systems: A Review". *Artificial Intelligence Review* 40(1):1–25.

Ristani, E., F. Solera, R. Zou, R. Cucchiara, and C. Tomasi. 2016. Performance Measures and a Data Set for Multi-Target, Multi-Camera Tracking. *In Lecture Notes in Computer Science* 9914 LNCS, 17–35. Cham: Springer.

Saroj, A., S. Roy, A. Guin, M. Hunter, and R. Fujimoto. 2018. "Smart City Real-Time Data-driven Transportation Simulation". In *Proceedings of the 2018 Winter Simulation Conference*, edited by M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 857–868. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Saurez, E., K. Hong, D. Lillethun, U. Ramachandran, and B. Ottenwälder. 2016. "Incremental Deployment and Migration of Geo-distributed Situation Awareness Applications in the Fog". In *Proceedings of the 10th ACM International Conference on Distributed and Event-based Systems* (DEBS'16), 258–269. New York, NY: ACM Press.

Terzis S. 2009. *Trust Management*. IEEE Computer Society Computing Now 8(9):1296–1296.

Voas, J., R. Kuhn, P. Laplante, and S. Applebaum, 2018. *Internet of Things (IoT) Trust Concerns*. NIST Cybersecurity White Paper, 17th October. https://csrc.nist.gov/CSRC/media/Publications/white-paper/2018/10/17/iot-trust-concerns/draft/documents/iot-trust-concerns-draft.pdf.

Xiong, L. and L. Liu. 2004. "Peertrust: Supporting Reputation-based Trust for Peer-to-Peer Electronic Communities". *IEEE Transactions on Knowledge and Data Engineering* 16(7):843–857.

Yan, Z., P. Zhang, and A. V. Vasilakos. 2014. "A Survey on Trust Management for Internet of Things". *Journal of Network and Computer Applications* 42:120-134.

Yigitoglu, E., L. Liu, M. Loper, and C. Pu. 2017. "Distributed Orchestration in Large-Scale IoT Systems". In *2017 IEEE International Congress on Internet of Things* (ICIOT), 22nd–25th October, Linz, Austria, pp. 58-65.

Yigitoglu, E., M. E. Gursoy, L. Liu, M. Loper, B. Bamba, and K. Lee. 2018. "PrivacyZone: A Novel Approach to Protecting Location Privacy of Mobile Users". In *2018 IEEE International Conference on Big Data*, 10th–13th December, Seattle, WA, 1238-1247.

Su, Z., L. Liu, M. Li, X. Fan, and Y. Zhou. 2015. "Reliable and Resilient Trust Management in Distributed Service Provision Networks". *ACM Transactions on the Web* 9(3), Article No. 14.

AUTHOR BIOGRAPHIES

MARGARET L. LOPER is the Chief Scientist for the Information & Communications Laboratory at the Georgia Tech Research Institute, an Associate Director for Trust for the GT Institute for Information Security and Privacy (IISP), and the Chief Technologist for the Georgia Tech Center for the Development and Application of Internet of Things Technologies (CDAIT). She has been involved in modeling and simulation research for more than thirty years, specifically focused on parallel and distributed systems. Margaret has taught simulation courses for both academic and professional education for ten years, and she has edited a book on how modeling and simulation are used in the systems engineering life cycle. Margaret is currently involved in projects related to Smart Cities, trust in Machine to Machine systems, and bringing modeling and simulation into K-12 education. She holds a Ph.D. in Computer Science from the Georgia Institute of Technology, a M.S. in Computer Engineering from the University of Central Florida, and a B.S. in Electrical Engineering from Clemson University. Her e-mail address is margaret.loper@gtri.gatech.edu.