COMPUTING edge

# INTERNET OF THINGS

**Also in this issue:**

- **Cybersecurity and the Future**
- **Mutual Dependence Demands Mutual Sharing**

JULY 2017

www.computer.org

IEEE

IEEE computer society

# PREPARE TO CONNECT

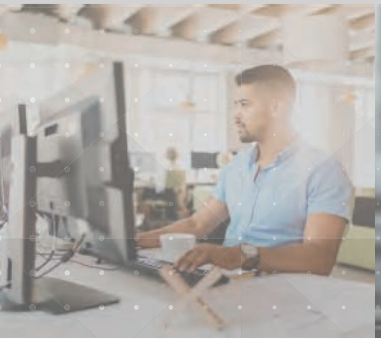
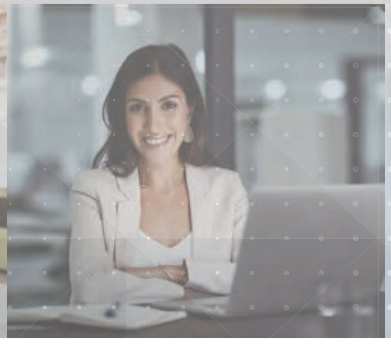
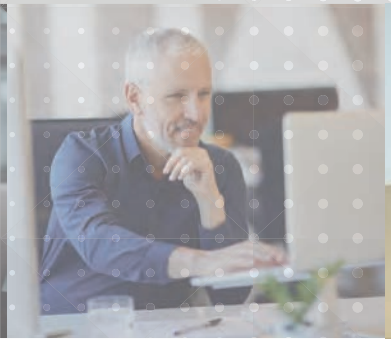
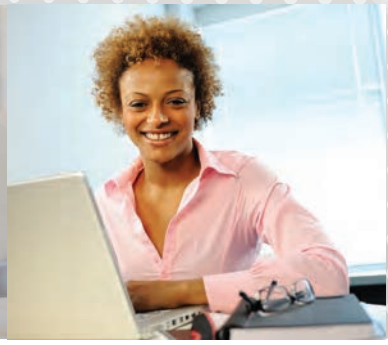

The IEEE Computer Society is launching **INTERFACE**, a new communication tool to help members engage, collaborate and stay current on CS activities. Use **INTERFACE** to learn about member accomplishments and find out how your peers are changing the world with technology.

We're putting our professional section and student branch chapters in the spotlight, sharing their recent activities and giving leaders a window into how chapters around the globe meet member expectations. Plus, **INTERFACE** will keep you informed on CS activities so you never miss a meeting, career development opportunity or important industry update.

**Launching this spring. Watch your email for its debut.**

IEEE computer society

**INTERFACE**

COMPUTING edge

IEEE COMPUTER SOCIETY http://computer.org • +1 714 821 8380

Printed with inks containing soy and/or vegetable oils

IEEE computer society

## STAFF



**Circulation:** ComputingEdge (ISSN 2469-7087) is published monthly by the IEEE Computer Society. IEEE Headquarters, Three Park Avenue, 17th Floor, New York, NY 10016-5997; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; voice +1 714 821 8380; fax +1 714 821 4010; IEEE Computer Society Headquarters, 2001 L Street NW, Suite 700, Washington, DC 20036.

**Postmaster:** Send address changes to ComputingEdge-IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Printed in USA.



## IEEE Computer Society Magazine Editors in Chief

***Computer***
Sumi Helal, *University of Florida*

***IEEE Software***
Diomidis Spinellis, *Athens University of Economics and Business*

***IEEE Internet Computing***
M. Brian Blake, *University of Miami*

***IT Professional***
San Murugesan, *BRITE Professional Services*

***IEEE Security & Privacy***
Ahmad-Reza Sadeghi, *Technical University of Darmstadt*

***IEEE Micro***
Lieven Eeckhout, *Ghent University*

***IEEE Computer Graphics and Applications***
L. Miguel Encarnação, *ACT, Inc.*

***IEEE Pervasive Computing***
Maria Ebling, *IBM T.J. Watson Research Center*

***Computing in Science & Engineering***
Jim X. Chen, *George Mason University*

***IEEE Intelligent Systems***
V.S. Subrahmanian, *University of Maryland*

***IEEE MultiMedia***
Yong Rui, *Lenovo Research and Technology*

***IEEE Annals of the History of Computing***
Nathan Ensmenger, *Indiana University Bloomington*

***IEEE Cloud Computing***
Mazin Yousif, *T-Systems International*

JULY 2017 • VOLUME 3, NUMBER 7

# COMPUTING edge


40

Internet of Things for Smart Cities: Interoperability and Open Data



## Departments

# Magazine Roundup

The IEEE Computer Society's lineup of 13 peer-reviewed technical magazines covers cutting-edge topics ranging from software design and computer graphics to Internet computing and security, from scientific applications and machine intelligence to cloud migration and microchip design. Here are highlights from recent issues.

## *Computer*

Competitions offer a compelling platform for engaging students and professionals in developing new technologies and skills. *Computer*'s July 2017 special issue on **challenge-based learning** explores various challenge-based approaches to education and technical innovation.

## *IEEE Software*

Today's abundance of processing power is affecting software engineering, according to "**How Abundance Changes Software Engineering**," from *IEEE Software's* May/June 2017 issue. By reducing the cost of failure, abundance changes how developers use computing technologies. Also, abundance changes developers' roles by moving their focus from the technology to its management.

## *IEEE Internet Computing*

Affordable computing and memory have enabled **innovative work in computer visualization and simulation**. As these technologies model high-resolution complexities, they offer new and sometimes unexpected answers but also generate new questions. This is the focus of "On the Simulation of Everything," from *IEEE Internet Computing*'s May/June 2017 issue, by Internet pioneer Vinton G. Cerf.

## *Computing in Science & Engineering*

The US Department of Energy's (DOE) **Exascale Computing Project** is a partnership between the DOE Office of Science and the US National Nuclear Security Administration. The project's mission is to transform today's high-performance computing (HPC) ecosystem via a multifaceted plan. The plan's elements include developing mission-critical applications of unprecedented complexity, supporting US national security initiatives, partnering with the US HPC industry to develop exascale computer architectures, and collaborating with US software vendors to develop an exascale-capable software stack suitable for industrial- and academic-scale systems. This is the focus of "The Exascale Computing Project," from *CiSE*'s May/June 2017 issue.

 Published by the IEEE Computer Society 2469-7087/17/$33.00 

*IEEE Security & Privacy*

Election security isn't just a matter of a secure system defending against attacks from an external adversary. It must also provide sound evidence of an accurate outcome. *IEEE S&P*'s May/June 2017 special issue on **electronic voting** contains articles on end-to-end verifiability, which can confirm that an election has handled votes correctly from casting to tallying.

*IEEE Cloud Computing*

One of the latest developments in cloud computing is called **serverless computing**, even though servers still handle the necessary processing. Serverless computing can save money for simple workflows. But first, users must model the economic impact of their architecture and operation choices, explains the author of "Be Wary of the Economics of 'Serverless' Cloud Computing," from *IEEE Cloud Computing*'s March/April 2017 issue.

*IEEE Computer Graphics and Applications*

**3D spatial user interface technologies** could make games more immersive and engaging. Although technologies such as stereoscopic 3D displays are now available for games, it's still unclear how their use affects play and performance. The authors of "Enhancing the Gaming Experience Using 3D Spatial User Interface Technologies," from *CG&A*'s May/June 2017 issue, discuss how these approaches affect gameplay.

*IEEE Intelligent Systems*

*IEEE Intelligent Systems*' May/June 2017 special issue addresses **computational advertising**. It includes articles on topics such as whether computational advertising is a paradigm shift for marketing, the use of verbal intent in semantic contextual advertising, and a three-phase approach for exploiting opinion mining in computational advertising.

*IEEE MultiMedia*

The author of "**Multimedia Research**: What Is the Right Approach?," from *IEEE MultiMedia*'s April–June 2017 issue, asks whether sufficient work is being done to figure out how to determine the best approaches to solving various types of research problems.

*IEEE Annals of the History of Computing*

**Micro Computer Machines**, established in Canada in 1971, was among the first companies to work on a personal microcomputer. The evolution of the company's views on software was representative of personal computing's many transformations in its early years, according to "MCM on Personal Software," from *IEEE Annals*' January–March 2017 issue.

*IEEE Pervasive Computing*

**Smart mobility technologies** help people access and exploit urban transportation options. New apps on the forefront of digitized transportation access will play a growing role in this process. The authors of "What Can We Learn from Smart Urban Mobility Technologies?," from *IEEE Pervasive Computing*'s April–June 2017 issue, note that these approaches not only affect mobility practices and user behavior but can also improve urban transportation planning.

*IT Professional*

*IT Pro*'s May/June 2017 special issue addresses the latest developments in **mobile data analytics**. This is an increasingly popular field that deals with analytics—especially big-data analytics—on resource-constrained mobile devices.

*IEEE Micro*

Designers who are making deep-learning computing more efficient can't rely solely on hardware. Incorporating software-optimization techniques such as model compression leads to significant power savings and performance improvement. This and related matters are discussed in "**Software-Hardware Codesign for Efficient Neural Network Acceleration**," from *IEEE Micro*'s March/April 2017 issue.

*Computing Now*

The Computing Now website (computingnow.computer.org) features **up-to-the-minute computing news** and blogs, along with articles ranging from peer-reviewed research to opinion pieces by industry leaders.

# A Haptic Compass for Navigation

**Lynette A. Jones,** MIT

*This installment highlighting the work published in IEEE Computer Society journals comes from* IEEE Transactions on Haptics.

**Figure 1.** Handheld haptic compass for use as a navigation aid.

Over the past decade, interest has grown in using haptic cues to aid navigation, both for people with visual impairments and for people with normal vision walking in unfamiliar environments. Many devices offer GPS-based audio instructions, but haptic cues are appealing because they're private and more practical in noisy environments. Numerous wearable devices based on vibrating motors have been developed and evaluated as navigation aids. In these systems, the location on the user's body at which the vibration occurs conveys information to the user about the intended direction of movement.

In "Development and Experimental Validation of a Haptic Compass based on Asymmetric Torque Stimuli," Jean-Philippe Choinière and Clément Gosselin describe the design, control, and experimental validation of a handheld haptic compass for use as a navigation aid (*IEEE Trans. Haptics*, vol. 10, no. 1, 2017, pp. 29–39). The device, shown in Figure 1, comprises a direct drive motor that generates open-loop torques around an axis orthogonal to the palm of the hand by accelerating the internal masses (the motor's rotor and flywheel) in one direction or the other. A haptic-signal feedback generator translates the user's location and orientation relative to an environmental target to determine the delivered torque's direction and amplitude. For example, if the user experiences a torque to the left, then the target is on the left; torque amplitude provides information regarding the distance to the target.

Experimental validation of the haptic compass indicated that with torque amplitudes greater than 10 mNm, users could accurately identify the initial direction in at least 90 percent of the test trials. Users responded faster to the torque signal as torque amplitude increased, averaging 8 seconds for torques greater than 50 mNm. Signals in the 5–15 Hz frequency range were perceived most clearly; the orientation of signals at lower frequencies took longer to be perceived. Users also performed best when the haptic feedback was proportional to the angular error between the user and the target.

In a more realistic evaluation of the haptic compass, the authors tasked blindfolded subjects with finding 15 indoor waypoints. Not only did all the users find the waypoints successfully, but they also considered the device intuitive to use with readily perceivable signals.

**LYNETTE A. JONES** is a senior research scientist in MIT's Department of Mechanical Engineering, and editor in chief of *IEEE Transactions on Haptics*. Contact her at ljones@mit.edu.

 Published by the IEEE Computer Society 

EDITOR'S NOTE

# Keeping Up with the Internet of Things

The Internet of Things (IoT) has enabled the design of new smart devices and has created new capabilities in areas such as information gathering; data analytics; automation; and the remote access of home, office, and industrial systems. July's *ComputingEdge* issue examines many of the IoT's exciting possibilities, as well as the challenges it faces.

Transforming IoT infrastructures into a general-purpose computing fabric could change how modern computation interfaces with our environment, according to *IEEE Software*'s "Software-Engineering the Internet of Things."

Early identification, mitigation, and prevention of problems with IoT applications could limit liability issues for system designers, notes *Computer*'s "The IoT Blame Game."

The authors of *IEEE Intelligent Systems*'s "IoT Quality Control for Data and Application Needs" discuss some of the challenges of and solutions to evaluating the quality of data in IoT systems.

"Continuous Authentication and Authorization for the Internet of Things," from *IEEE Internet Computing*, addresses the development of these security techniques for the IoT.

*IEEE Micro*'s "Visual IoT: Architectural Challenges and Opportunities; Toward a Self-Learning and Energy-Neutral IoT," includes invited position papers about these two potentially important IoT approaches.

"Osmotic Flow: Osmotic Computing + IoT Workflow," from *IEEE Cloud Computing*, focuses on a new way to efficiently execute IoT services and applications at the network edge, to help cope with the volume and variety of big data that IoT devices produce.

The authors of *IEEE Internet Computing*'s "Internet of Things for Smart Cities: Interoperability and Open Data," provide a case study of Sweden's GreenIoT platform.

"IoT: From Sports to Fashion and Everything In-Between," from *IEEE Pervasive Computing*, looks at how pervasive computing and the IoT now affect many diverse fields.

This issue also includes articles on topics other than the IoT:

- *Computer*'s "Cybersecurity and the Future" discusses the challenges security will have to meet as technology changes in the coming years.
- As news migrates to mobile phones, media companies are turning to data visualization to whet readers' appetites for stories they can read at length on their home or work computers, according to "The Need to Help Journalists with Data and Information Visualization," from *IEEE Computer Graphics and Applications*.
- To benefit fully from today's technology, we must complement it with an inexpensive system that protects confidentiality and reports on the volume, pattern, and character of incoming digital attacks, contend the authors of *IEEE Security & Privacy*'s "Mutual Dependence Demands Mutual Sharing."

Editor in Chief: **Diomidis Spinellis**
Athens University of Economics and Business, dds@computer.org

# Software-Engineering the Internet of Things

Diomidis Spinellis

**ENIAC WAS BUILT** during the Second World War, from 1943 to 1945. Many consider it the first electronic, general-purpose, large-scale digital computer. Picture it as a room encompassing 36 racks, three printer panels, a card reader, and a card punch. Each rack used hundreds of vacuum tubes to perform a specific function. Many racks acted as accumulators: they received pulses corresponding to the digits of a decimal number and increased accordingly the number stored in them. Others were more specialized; a 1944 floor plan has racks labeled multiplier, partial product, square rooter, denominator, multiplicand, and so on. Three function table racks, initialized through switches, could be carted around on wheels.

Today we build, connect, and configure most Internet of Things (IoT) systems by linking together their sensor, actuator, and computing nodes through cloud infrastructures, mobile apps, and the sharing of security credentials. Similarly, ENIAC was programmed by setting up function tables and switches and connecting its units in the manner required for solving a particular problem, such as generating sine and cosine tables or computing artillery trajectories and shock wave reflections. As you can imagine, such programming was time-consuming and error-prone.

Then, in 1948 a remarkable thing happened. Inspired by the design of EDVAC (Electronic Discrete Variable Automatic Computer), discussed over summer school lectures at the University of Pennsylvania's Moore School, ENIAC's designers realized they could wire it in a revolutionary way. The wiring wouldn't solve a particular numeric program. Instead, the designers would repurpose some of ENIAC's accumulators so that it would read instructions prescribing what actions to perform from its numeric function tables. Think of this as building a command interpreter by assembling together already existing discrete electronic components. Thus, the new wiring transformed ENIAC into a versatile stored-program computer. Rewiring IoT infrastructures into a general-purpose computing fabric can similarly change how modern computation interfaces with our environment.

## A Maze of Problems ... and Ways Out

No doubt, a paradigm shift from balkanized IoT applications to an integrated infrastructure in which individual IoT nodes are first-class citizens raises formidable challenges. Start with requirements. In a standalone IoT application, it can be easy to satisfy a major functional requirement—say, home security—by controlling the balance of diverse nonfunctional requirements, such as performance, reliability, and usability. However, when multiple IoT nodes and applications get integrated, diverse requirements will interfere with each other (what happens when a burglar triggers a fire alarm?), requiring difficult prioritizations and multicriteria decision making. Given the fluid nature of IoT systems, many decisions we make today during requirements' elicitation might need to be obtained dynamically as the systems operate.

*IEEE Software* Mission Statement

To be the best source of reliable, useful, peer-reviewed information for leading software practitioners—the developers and managers who want to keep up with rapid technology change.

 Published by the IEEE Computer Society 2469-7087/17/$33.00 

Designing and constructing integrated IoT systems will also be tough. Standardization in the IoT area is still in its infancy, requiring complex adaptation layers. Requirements associated with control, privacy, and reliability might dictate design decisions that will be at odds with IoT nodes' processing capacity, power budget, bandwidth, and ubiquity. The rewired ENIAC's 60 "order codes," which we today would call its instruction set, required two decades of innovations in programming languages, operating systems, and databases to provide us with the tools and abstractions we now routinely use to build computing applications. Progress of a similar ambition and scale might well be required to truly harness the IoT.

IoT nodes' long-term maintenance will be another nightmare. Many devices will be embedded into buildings, streetlights, bridges, cars, appliances, and other places with lifetimes at least an order longer than the typical PC, smartphone, or server. The IoT nodes will require regular corrective and perfective maintenance, even as their vendors inevitably switch priorities or go out of business over time. We must therefore come up with ways to smooth the handover of IoT devices between vendors.

Also, maintaining a system made up of hundreds (or millions) of diverse embedded IoT nodes is a completely different ballgame than looking after a monolithic cloud application and its systems software. By now, we have some limited experience in such tasks through component-based ecosystems, such as Node.js and microservice architectures. With IoT maintenance, we must apply and extend this knowledge to devices, longer timescales, and a much larger scope.

Then, consider configuration management, which will entail accommodating diverse stakeholders and dynamically changing systems. In DevOps settings, we're already facing tricky problems when a vendor's configuration clashes with locally implemented changes. This problem will only multiply when systems with multiple vendors evolve over time with configurations being set up by unspecialized and untrained users. Unless the state of the art improves dramatically, we might face the choice between unrealized IoT promises and a mess of conflicting, incompatible configurations that will make the 1990s DLL (dynamic linked library) hell appear like heavenly peace. On the other hand, the capabilities of modern decentralized configuration management systems, container technologies, and package managers might offer us the building blocks of a possible solution.

Nailing down the quality of IoT-based systems will also be an arduous, long-term task. Start with security. By definition, IoT systems will be interconnected and offer access to the physical world—a dream come true for cyberwarriors, spooks, and cyberterrorists. Today, we're witnessing the edifice's cracks as rogues take over IoT devices to launch distributed denial-of-service attacks; tomorrow's attacks might be deadly.

Complaints from people unable to control their IoT-enabled fridges hint that usability will be another potent source of ridicule and problems. Furthermore, the physical-world interfaces of IoT nodes mean that reliability failures, of which current software has too many, can have much more serious repercussions than the inconvenience of an odd lost file or application crash.

## A CHANGING OF THE GUARD

After four years of dedicated service, Ipek Ozkaya handed the baton of the *IEEE Software* advisory board chair to Rafael Prikladnicki, whom our readers already know as the editor of the Voice of Evidence column. Ipek will continue to serve as associate editor of the magazine's departments, taking over from Richard Selby, who staffed this post over the past two years. Also, after delivering two years of phenomenal growth in *IEEE Software*'s social media presence, Alexander Serebrenik passed on the corresponding management responsibility to his team member Damian Andrew Tamburri.

*IEEE Software* advertises new volunteer openings through social media and the SEWORLD mailing list (www.sigsoft.org/resources/seworld.html). Keep an eye out if you want to join our team!

Finally, cooperating to achieve the best overall performance of an IoT system, when each device might be selfishly trying to maximize its own, will also be tricky. We've successfully solved a similar problem in the case of Internet bandwidth allocation, through the design of the TCP protocol. Similarly ingenious approaches might well be required for the IoT.

Inevitably, the problems I outlined will feed into an engineering-management challenge: coordinating multiple stakeholders with conflicting interests. One option might involve restricting the setup of IoT systems to walled gardens with strictly defined standards, processes, and compliance testing. This approach might work in a simple controlled environment, such as a small-business building or plant. However, it will rob us of the innovation that open environments can spur and will likely severely limit IoT applications and their impact.

So, unless a winner-takes-all scenario materializes, a closed system will be unsuitable for environments with multiple stakeholders, such as cities, private residences, large office buildings, and mobility applications. Instead, market-based mechanisms should probably be introduced as a way to reduce the friction between technology, social interactions, and the physical world.

Blighted by budget and schedule overruns followed by unreliable operation, ENIAC's birth was anything but auspicious. However, thanks to the perseverance, ingenuity, and openness of the people behind it, it became a defining milestone for modern computing. With hard work and some luck, the IoT can usher in similar changes to how we interact with the physical world.

*This article originally appeared in* IEEE Software, *vol. 34, no. 1, 2017.*

EDITOR **ROY WANT**
Google; roywant@google.com

# The IoT Blame Game

**Jeffrey Voas,** IEEE Fellow
**Phillip A. Laplante,** Pennsylvania State University

*Applications in the Internet of Things offer many benefits and risks. Poor design, unplanned interconnections, and human adversaries raise the stakes and liabilities for system designers. Early identification, mitigation, and prevention of these threats can help limit liability issues.*

## FROM THE EDITOR

This edition of the IoT Connection considers the legal ramifications of faulty Internet of Things (IoT) devices, their associated infrastructure, and the heterogeneous patchwork of systems that will become the global IoT. As we begin to rely on this ubiquitous network for critical tasks, a fault might result in serious consequences. The authors discuss why the IoT will challenge our legal system to do the right thing when lawsuits arise. —*Roy Want*

The Internet of Things (IoT) is expected to drive a technological revolution because the easy connection of and communication among "things" will allow engineers to quickly build large numbers of intelligent networks and associated infrastructure. The IoT is expected to impact everything from food production and healthcare to transportation, communication systems, and most forms of automation. In short, the IoT will impact everyone in many ways.

With such great societal impact, it's imperative that IoT-based systems are trustworthy. Here, "trustworthy" means that a system should be secure and reliable, and it should possess many other attributes associated with quality.[1,2] In addition, privacy in such systems is of particular importance because they'll likely generate large amounts of data due to their sensing and monitoring capabilities.[3] Therefore, techniques, tools, and methods to mitigate diverse trust challenges are needed before such systems can safely manage daily life.

IoT systems built from "things" are increasing in diversity, scale, and number. But as they move from laboratory proofs of concept to fielded applications, new and unforeseen risks are likely to emerge. Possibly the most serious of these issues surrounds cyber-physical systems that control life-critical applications.

Liability matters quickly manifest because of the potential for unplanned interactions between critical and noncritical things. Further, the heterogeneity and lack of ownership and control of many of the things in a specifically purposed network of things (NoT)[3] can exacerbate the problem. Ultimately, parties involved in a failed

2469-7087/17/$33.00 © 2017 IEEE Published by the IEEE Computer Society 

system will start to assign blame, informally, then move on to public pleadings, and, ultimately, the matter will be dealt with in legal proceedings—we call this the "blame game." The blame game and liability go hand in hand.

Because conducting a root-cause analysis of failure in these kinds of systems is complex, we explore some of the more interesting challenges and how they relate to the matter of liability. We also dismiss the common belief that standards, for any aspect of a system composed of things, can take a one-size-fits-all approach, because "all" can't be well defined or bounded.

> It can be very difficult to properly attribute blame when a system is in contact with hundreds of leased third-party products and services and their interactions.

## SYSTEMS AND VULNERABILITIES

We consider critical systems to be those involving infrastructure that touches or supports human life in a way that could cause harm in the event of malfunction. A short list of critical systems includes those in

- telecommunications,
- the water supply,
- electrical power generation and distribution,
- road transportation,
- railway transportation,
- air transportation,
- public safety services, and
- healthcare.[4]

In any of these applications, certain types of failures could lead to serious injury or fatality.

Applications that rely on networks of connected things will receive inputs from sensors, cameras, various other devices, and even social media feeds. These inputs can be transmitted over wired and wireless communications to an aggregation point, where processors apply various algorithms to create output signals, which are then sent to other receptors that facilitate decisions or actuators. Each one of these interactions presents a failure point that could result in catastrophe.

For example, consider an application that provides seamless tracking of a container of critical medical supplies (or even transplant organs) from a large hospital to a rural hospital far away. The container leaves the large hospital and is transferred to an ambulance, which travels the roadways to an airport. The container is then transferred onto an airplane, and then, upon landing, to a train, then to another ambulance, and it finally arrives at the rural hospital. During its travels, the container has to cross through many different IoT ecosystems, complying with different standards (highway, airport, railway, hospital) and performing handoffs through a variety of communications links (including in the air, where communication is usually forbidden). Imagine the engineering problems in building this application such that patients' lives aren't put at risk, and also consider the number of things that affect the application that might be difficult to trust.

There's another problem with these complex, multi-domain systems. Connectivity through handheld devices, smart homes, smart cars, and wireless-enabled devices increases the attack surface for hackers to exploit. In fact, vulnerabilities in critical infrastructure systems have already been reported. For example, there are known vulnerabilities in certain smart power meters, and in older supervisory control and data acquisition (SCADA) systems that control many of the critical infrastructure systems around the world.

SCADA's advantage—that it can remotely control large operations from great distances, such as an offshore oil facility—is also its downside: the communications that enable remote control and monitoring can be compromised. More recently, certain traffic-control sensors and network-connected smart LED lightbulbs were shown to be vulnerable to hacking. Such vulnerabilities can be found in any of the ecosystems in the aforementioned container transport example: automotive systems, railway transportation systems, and even in-flight systems. In fact, vulnerabilities can be found and exploited anywhere, including in-home appliances.[4]

## IDENTIFYING THE SUSPECTS

To mitigate this problem, IoT app developers must first answer the following question: Will we own and control all of the assets (things, clouds, data, sources of data, communication channels, software, and so on) that comprise our NoT? If not, we're in a compromised position from the outset. If we answered yes, we're in better shape than most.

Let's look at a short list of the risks we face by answering no.

- *Leased data*. This is data that comes from suppliers at the time of their choosing and with the integrity of their choosing. We can try to SLA (service-level agreement) our way into a better agreement here, but good luck enforcing it. This is an immediate supply-chain concern. It's also a business concern—what if our competitors lease the same data from the same supplier? Are all being treated fairly?
- *Faulty interfaces*. Interfaces are at the endpoints of the "veins" and "arteries" of an NoT. We'll

likely rely on a wireless service provider—is that provider trustworthy?
- *Defective things*. This aspect includes every third-party thing. If we didn't create that thing, we're subjected to whatever the third party's limited warranty says. And it might not be in our favor or even useful.
- *Faulty or subpar architecture*. Did we architect your network, or did we contract that task out? Were the best things used in implementing the architecture? Were security and privacy even considered in that architecture? Did we have enough time to do an economic tradeoff analysis for when to select more expensive things for higher integrity and when to use lower-quality components? We can easily over-engineer a solution here.
- *Data tampering and integrity*. Data is the "blood" that flows through a network. How secure is our data from malicious tampering, delay, or theft?
- *Expected operational usage*. Do we have a good idea of the environment that our NoT will operate in? Did we design our network for that or was this task hired out? Getting this wrong will almost certainly cause problems from the outset of deployment.

This short list gives a sense of how difficult it can be to properly attribute blame when a system is in contact with hundreds of leased third-party products and services and their interactions.

## NEED A LAWYER?

Liability frameworks for the aforementioned issues (and others) should be jointly created by governments, industry, risk analysts, and insurers.[5] But that approach is based on a gross assumption: that these things and the networks they populate will have lifespans long enough to sort out the legal challenges that might arise when something goes wrong. Leased data from a third-party supplier might instantly enter into a network in near-zero time, and then the next set of data enters similarly. Who's keeping track of these near-instantaneous data transactions? If no one, no attribution of blame is possible. The point here is that there's a certain disposability attribute to the things and services that might be used, temporarily, in an NoT, making attributions of liability difficult and not timely enough to pursue.

Liability in an NoT is characterized by two questions:

- Who's liable when a thing fails?
- What's the probability that a problem event will occur? In other words, what is the risk associated with using a thing?[6]

For systems of things, we can amend these questions to

- Who's liable when the system fails?
- What's the possibility of system failure?

Let's start with the first question. For a homegrown, non-interconnected system, the answer has to be the developer. But for systems that are connected to other systems locally and through the Internet, the answer becomes more difficult. According to a lawyer specializing in the IoT and Internet liability: "In case of (planned) interconnected technologies, when there is a 'malfunctioning thing' it is difficult to determine the perimeter of the liability of each supplier. The issue is even more complex for artificial intelligence systems involving a massive amount of collected data so that it might be quite hard to determine the reason why the system made a specific decision at a specific time."[7]

The answer to the second question is very difficult to determine. A powerful technique for determining the risks of a system-level failure would involve "fault injection dynamic methods that simulate the effects that real faults will most likely have as opposed to simulating the faults themselves as a means of quantifying the risks created by the software component of a system."[6] But until we can accurately and scientifically measure these risks, we likely won't have a means for probabilistically and mathematically bounding and quantifying liability.

## VETTING AND CERTIFICATION

In a previous work, Jeffrey Voas described three generic areas of reducing the risk of failures (and hence liability): process, personnel, and product certification.[8]

**People.** Although never fail-proof, the risk of failure for any kind of engineered product can be significantly reduced by having the best people work on it. High-quality engineers can be characterized by their educational background and experience. Some have also argued that for safety-critical systems in the IoT era, professional licensure is an important consideration.[4] Other ways of ensuring excellence in engineering personnel include certifications, ongoing training, and observance of professional codes of ethics. IEEE and ACM are updating their codes of ethics so that they're appropriate for the complexities of IoT systems. The problem here is that we don't have a consistent understanding of what makes a computer scientist or other engineering graduate an "IoT engineer." Therefore, certifying IoT engineers is currently impossible.

**Process.** Reducing the risk of failure requires a comprehensive approach involving selecting the correct engineering lifecycle model and associated processes, embracing process discipline and standards compliance, understanding complex interactions, and selecting the right tools for use throughout the system's lifecycle. These activities

describe those used by the highest-level software professionals. In fact, we should insist on appropriate certification and licensure for electrical, computer, control systems, and software engineers who build critical infrastructure systems. Depending on the domain (power, nuclear, or civil), other licensed engineers will also need to be involved. However, presuming that a great process results in a great product is overly simplistic and very risky.

**Product.** The third way to reduce the risk of failure is through product certification. This approach assesses the things within the IoT to determine their quality. Here, rigorous, repeatable, and reproducible assessment technologies are needed to establish a thing's trustworthiness. If inferior assessment technologies are employed, we might quickly return to "the quagmire created by process and personnel certification: not knowing how good or bad the behavior of the software will be."[9] Product certification has long been the holy grail of certification—and it remains elusive, except in cases where time-to-market and costs aren't as important (for example, safety-critical systems under government regulation).

Current certification approaches to things within the IoT are, at best, mired in quicksand. We can expend a lot of time and effort on certification with questionable benefits. This situation is no different than the difficulties encountered in vetting mobile apps or any other component that is a miniscule part of total system functionality. These three Ps outline the basics of the blame game: we can accuse the people, products, or process, or some combination of the three.

## CHAIN OF CUSTODY

Interactions (both planned and unplanned) between critical and noncritical systems create significant problems of risk and liability. These interacting, dynamic, cross-domain ecosystems create the potential for increased threat vectors, new vulnerabilities, and risks.

Consider, for example, this scenario: hacked refrigerator software interacts with an app on a woman's smartphone, installing a security exploit that can be propagated to other applications with which her phone interacts. The woman enters her automobile and her phone interacts with the vehicle's operator interface software, which downloads the new software, including the defect. Unfortunately, the software defect causes an interaction problem (for example, a deadlock) that leads to a failure in the software-controlled safety system during a crash, leading to injury. The potential for this chain of events to play out demonstrates why interoperability in IoT technologies is so challenging regarding identifying and mitigating risk, and assigning blame when something goes wrong.

Another way in which interactions in the IoT network introduces murky legal issues is when different components in the application reside in locations with different jurisdictional law with respect to privacy and data ownership. This heterogeneous, grand-scale, distributed nature of IoT-based systems will require a corresponding legal framework that can adequately take into account scalability, verticality, ubiquity, and interoperability.[9]

## KEEPING WATCH

Most consumers will hope that the IoT offers reasonable levels of security, reliability, and privacy. But how can these assurances be given without prescriptive architectures and certified or verified things?

Possibly the best approach for consumers would be to securely audit and log their internal and external operations and data interactions. The presence of an auditing system that can operate independent of any IoT vendor of third-party things will foster increased vendor interoperability as well as an acceptance of standards among IoT technologies. Another advantage of auditing and logging is that these processes offer the ability to increase reliability and resilience without requiring major changes to the architectures of specifically purposed NoTs. Ultimately, through increased auditing, IoT systems will continue to improve in terms of both security and reliability. End users will benefit from improved operational transparency, empowering them to identify components that can be used together, thus improving IoT systems' utility.

Of course, we don't know if the average consumer can understand the results from such an auditing and logging process. Perhaps the goal of this activity would be more akin to creating a dashboard with the most relevant and mitigatable information from a variety of tools. Although consumers aren't all computer scientists or electrical engineers, they do understand the importance of antivirus software. They buy and install it, and they're comfortable doing so. Ideally, the same could happen with IoT auditing and logging.

The owner of an NoT will have to take some level of responsibility of the risks of the things and services in the system. Fortunately, there are a few commercial products available that attempt to help consumers limit liability for their NoTs, including Dojo Labs (www.dojo-labs.com), Bitdefender BOX (www.bitdefender.com/box), CUJO (www.getcujo.com), Norton Core (us.norton.com/core), and Internet of Things Armor (www.iotarmor.net). We hope this marketplace continues to grow.

## REFERENCES

1. J. Voas and G. Hurlburt, "Third Party Software's Trust Quagmire," *Computer*, vol. 48, no. 12, 2015, pp. 80–87.
2. C. Kolias et al., "Learning Internet of Things Security 'Hands-On,'" *IEEE Security & Privacy*, vol. 14, no. 1, 2016, pp. 37–46.

3. J. Voas, *Networks of 'Things,'* NIST Special Publication 800-183, July 2016; nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-183.pdf.
4. P.A. Laplante, "IEEE Insight, Critical Infrastructure and the IoT: A Licensing Perspective," *IEEE-USA InSight*, 18 Sept. 2015; insight.ieeeusa.org/insight/content/views/175444.
5. M. Chui, M. Löffler, and R. Roberts, "The Internet of Things," McKinsey Q., March 2010; www.mckinsey.com/industries/high-tech/our-insights/the-internet-of-things.
6. J. Voas et al., "A 'Crystal Ball' for Software Liability," *Computer*, vol. 30, no. 6, 1997, pp. 29–36
7. G. Coraggio, "The Internet of Things and its Legal Dilemmas," VC Experts Blog, 15 Dec. 2016; blog.vcexperts.com/2016/12/15/the-internet-of-things-and-its-legal-dilemmas.
8. J.M. Voas, "Limited Software Warranties," *Proc. 7th IEEE Int'l Conf. Workshop Eng. Computer-Based Systems* (ECBS 00), 2000; doi: 10.1109/ECBS.2000.839861.
9. Z. Yan, P.Z. Zheng, and A.V. Athanasios, "A Survey on Trust Management for Internet of Things," *J. Network and Computer Applications*, vol. 42, 2014, pp. 120–134.



**JEFFREY VOAS** is an IEEE Fellow. Contact him at j.voas@ieee.org.

**PHILLIP A. LAPLANTE** is a professor of software and systems engineering at Pennsylvania State University and an IEEE Fellow. Contact him at pal11@psu.edu.



*This article originally appeared in* Computer, *vol. 50, no. 6, 2017.*



Editor: **Amit Sheth,** Kno.e.sis at Wright State University, amit@knoesis.org

# IoT Quality Control for Data and Application Needs

**Tanvi Banerjee and Amit Sheth,** *Kno.e.sis Center at Wright State University*

With the rapid growth of sensors and devices that communicate—that is, the Internet of Things (IoT)—smart devices have permeated every facet of modern life. These IoT devices are within our bodies, on our bodies, in the environment both inside and outside our homes, observing our behavior patterns on a day-to-day basis, and assisting in production systems and surveillance. Figure 1 highlights some of the more popular IoT applications in the world.

However, with these sensors' ubiquity and pervasiveness comes vast amounts of data that need to be processed and analyzed to extract meaningful or actionable information from the data for recommending appropriate changes in the real world. This requires using not only semantic approaches,[1] but also data streamlining to ensure that the decisions made are not erroneous. Moreover, due to the sheer volume of the data from these IoT devices, any errors from user entry, data corruption, data accumulation, data integration, or data processing can snowball, causing massive errors that can detrimentally affect the decision-making process. Consequently, there needs to be a clear understanding of the challenges associated with data quality and a way to evaluate and ensure that data quality is maintained for different applications.

## Mapping the OSI Framework to IoT Quality

There are several implementation concepts behind determining the IoT architectures that are application-driven and face different challenges, whether at the hardware level, software level, or integration. Specifically from the Open Systems Interconnection (OSI) perspective, three quality areas need to be examined: the device level (data link layer), the network level (network layer), and the application level (presentation and application layers). The choice of IoT devices and design protocols in the IoT systems determine the choice of implementation design, as well as analysis algorithms to achieve the optimal quality of service.[2] Specifically, the lower OSI layers have been extensively investigated in several studies to extract and transmit the raw sensor data from IoT devices through protocols such as MAC or IEEE 802.3.[2] However, the higher layers often get overlooked, especially from the perspective of the target domain. Using two scenarios, we will show how data quality can be evaluated contextually and how the different OSI layers are affected by the specific user needs of the system.

Before we get into the use cases, we need to understand data quality, "the degree to which a set of inherent characteristics fulfills the requirements."[3] Within quality, we have two categories: specification and conformance quality.[4] Specification quality refers to how well a device matches with other similar devices in that domain. Conformance quality looks at the "correctness" or the veracity of the readings from the device. We add one more quality control that needs to be examined in this setting: the device's semantic quality. Clearly, the interpretation of the data from the sensor has a key role to play in an application targeted toward a specific healthcare requirement, as our two use cases show.

## Use Case 1: Lung Function for the Elderly

With the rapid growth of wearable technologies in the mobile industry, the healthcare industry is pushing the boundaries of continuous activity monitoring using wearables.[5] In this scenario, we consider the use of a popular wearable vest called Hexoskin for measuring physiological changes in older adults.[6] Specifically, let's look at one of the physiological sensors, the lung volume measures that compute the tidal volume of the lungs using the last inspiration (80 mL to 10 L) at a frequency of 1 Hz, and the frequency of the inspiration and expiration events.

 Published by the IEEE Computer Society 

Now let us look at some of the challenges in the OSI layers that need to be considered for this use case.

## Applying Data Semantics in the Presentation Layer

In terms of preprocessing the sensor data on the Hexoskin, the vendor's website (www.hexoskin.com) mentions baseline change and noise detection but does not provide details into what this entails. In particular, one aspect of these wearables that is overlooked when it comes to its use in the nonactive population is that the baseline or even the noise measurement can differ significantly in older populations. In fact, articles[7] and studies[8] describe the lung change in the older population as similar to emphysema. This can lead to a low correlation in lung function and activity, and must be considered when vital sign functions are interpreted.[9] Furthermore, this can affect the analysis when using the vest on the elderly population. This constraint is not surprising given the fact that the sensor was primarily built for fitness measurement. Figure 2 shows how understanding the challenges at the upper layers of the OSI model can improve the performance of wearable systems in healthcare applications.

Consider an older man, Bob, who is wearing the Hexoskin vest for activity measurement. Suppose the lung function readings (data) for him are FEV1 (forced exhaled volume in 1 second) = 2.04 L and FVC (forced vital capacity) = 3 L. Although the values themselves have no significance for chronic obstructive pulmonary disease (COPD), the FEV1/FVC ratio (also called the Tiffeneau-Pinelli index) represents the information as the proportion of a person's vital capacity that he or she can expire in the first second of forced expiration. We see from the FEV1/FVC ratio that this index is greater than 0.7. This is where the knowledge indicates that there may be concern for obstructive pulmonary disease such as emphysema.[10] However, this could also be a manifestation of aging,[8] so we need to dig one step further in the process and ask contextual questions, such as whether Bob has a history of smoking, or conduct additional tests, such as a flow volume loop, to diagnose whether Bob has a lung condition. If in fact the additional tests indicate that Bob has emphysema, the wisdom (relevant actionable medical science) comes in its treatment and handling the day-to-day variabilities in the symptoms that allow Bob an improved quality of life. In this scenario, the data and information part of the pyramid map to the presentation layer, which checks the context and veracity of the data readings, wherein the data are interpreted to enable semantic quality control of the IoT device readings for this application. Similarly, as Figure 2 shows, the knowledge and wisdom portions of the pyramid map to the application layer to incorporate additional data sources for more meaningful data analysis.

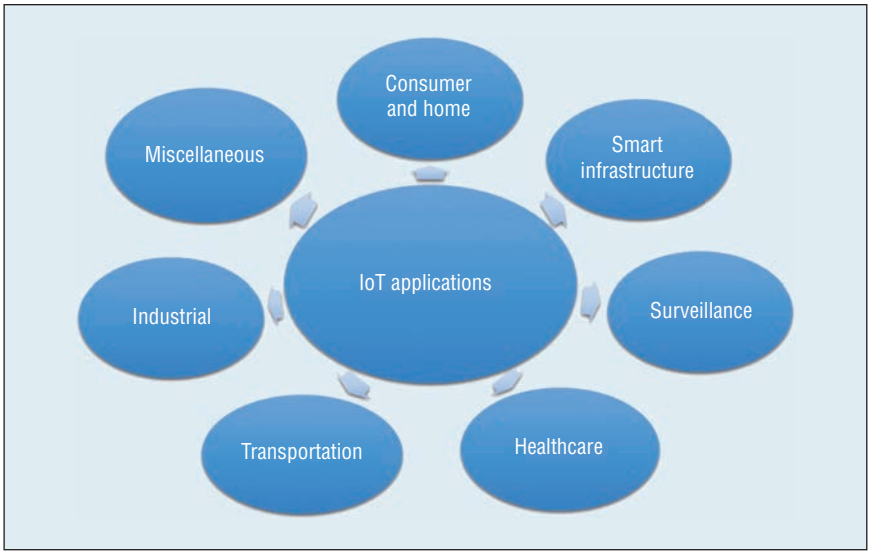


**Figure 1. Key application areas using the Internet of Things (IoT). This article focuses on the healthcare applications of the IoT.**

## Incorporating Knowledge at the Application Layer

One way to overcome this manifestation of aging in the sensor readings is to use the raw sensor information and change the baselines for the older population or for populations with certain health conditions like emphysema. Here, the semantic mapping between the sensor values and knowledge of the population using the system needs to be taken into account to ensure that the results are accurate and useful. To potentially incorporate this information, we use ontologies or knowledge representations to annotate the data specifically for this application. An example of such an ontology is the semantic sensor network ontology.[11] Creating such shared semantic definitions helps integrate new data into historical, temporal, and spatial contexts. Definitions of sensors and their capabilities are also useful for quality reasoning. For example, if the accuracy of a sensor depends on phenomena other than that which it measures, then a specification of this can be used as a guide to search for spatially and temporally related measurements of the phenomena on which the accuracy depends, which then defines the application's quality metrics.

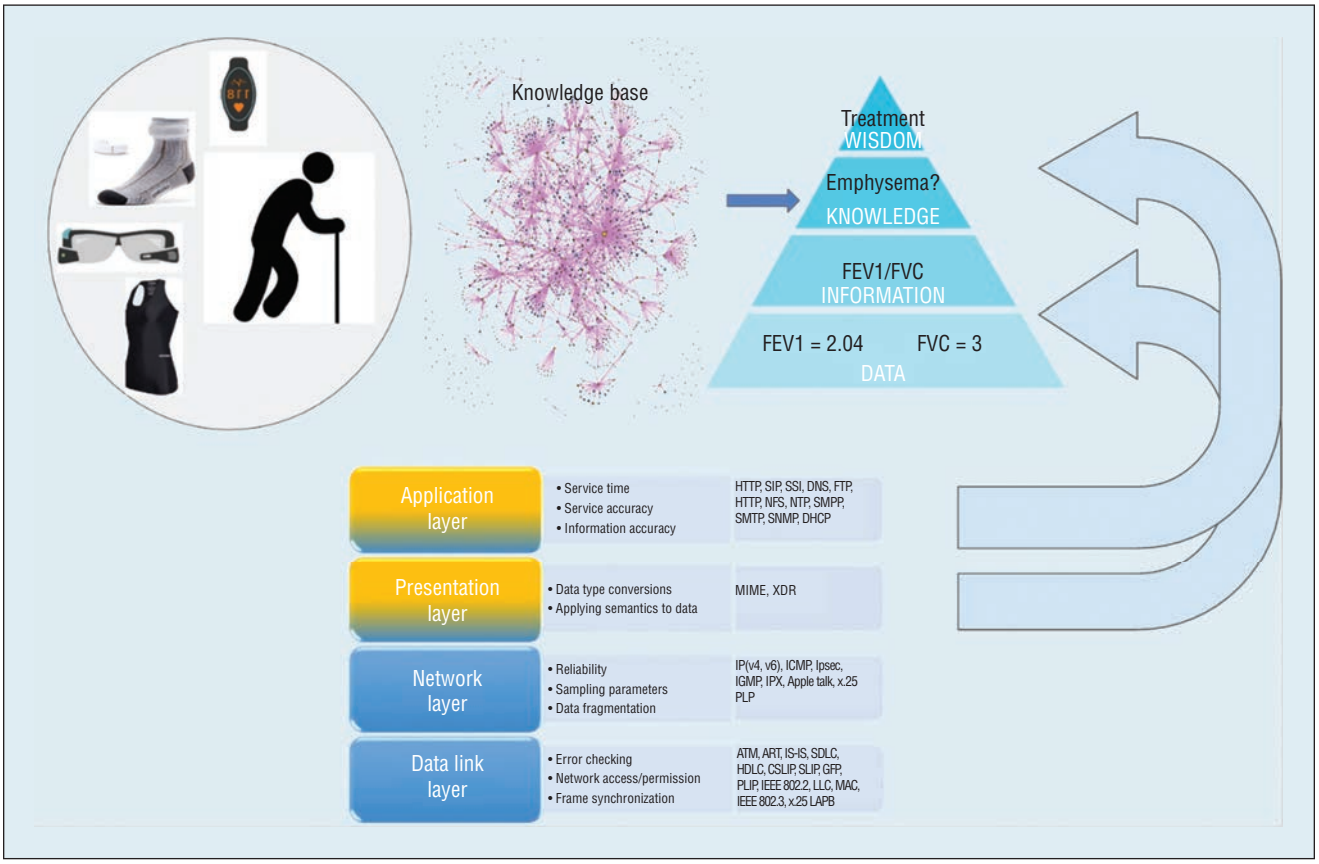


**Figure 2. The bottom of the figure shows four of the OSI layers: data link, network, presentation, and application. The top shows the DIKW (data-information-knowledge-wisdom) pyramid that maps wearable sensor data to meaningful information. The current use case is in the area of lung function analysis in the elderly population.**

Specifically, in our example from Figure 2, for the older adult population, we can restrict the operating range (property `ssn#MeasurementRange`) to accommodate the differences in the breathing measurements for the target population, which can help improve the accuracy of the readings through identification of more errors or discrepancies in the data values. Many of the other parameters or entities described in the ontology are specifically left undefined to fit user requirements and enable reusability, which can be leveraged for specific target populations to improve user-centric applications, especially in healthcare.

## Use Case 2: Smart Home System

In this use case, consider a smart home-monitoring system that uses several IoT devices to measure the physical and cognitive health conditions of older adults living independently in their homes. Such a system can be used to detect falls and to study the residents' continuous behavior pattern and generate alerts when the residents deviate from their normal behavior. Apart from the wearable sensors discussed earlier that can measure physiological changes in the residents, environmental sensors are also placed in this smart environment. These include a depth-sensor-based system that is in place for continuous anonymized fall monitoring and activity analysis, including sedentary behavior.[12,13] Moreover, wireless motion sensors are in place in the environment to study the interactions of the people in the environment, as well as to further examine their behavior patterns, such as measuring the time away from home[14] and movement within the living spaces.[13,14] Figure 3 shows such a system in the home of an older couple; the IoT devices are all routed through a common channel via the Internet and stored in a secure database. Corresponding behavior analysis is shared with the clinician as well as the couple's family members. In this use case, we will discuss two challenges in terms of the network and application layers that relate to a more complex multimodal data fusion IoT-focused application.

### Incorporating Data Integration at the Network Layer

The fall-detection system we described earlier requires a common framework that can integrate wearable sensors with environmental sensors such as the wireless body/personal area network (WBAN).[15] WBAN allows long-term,

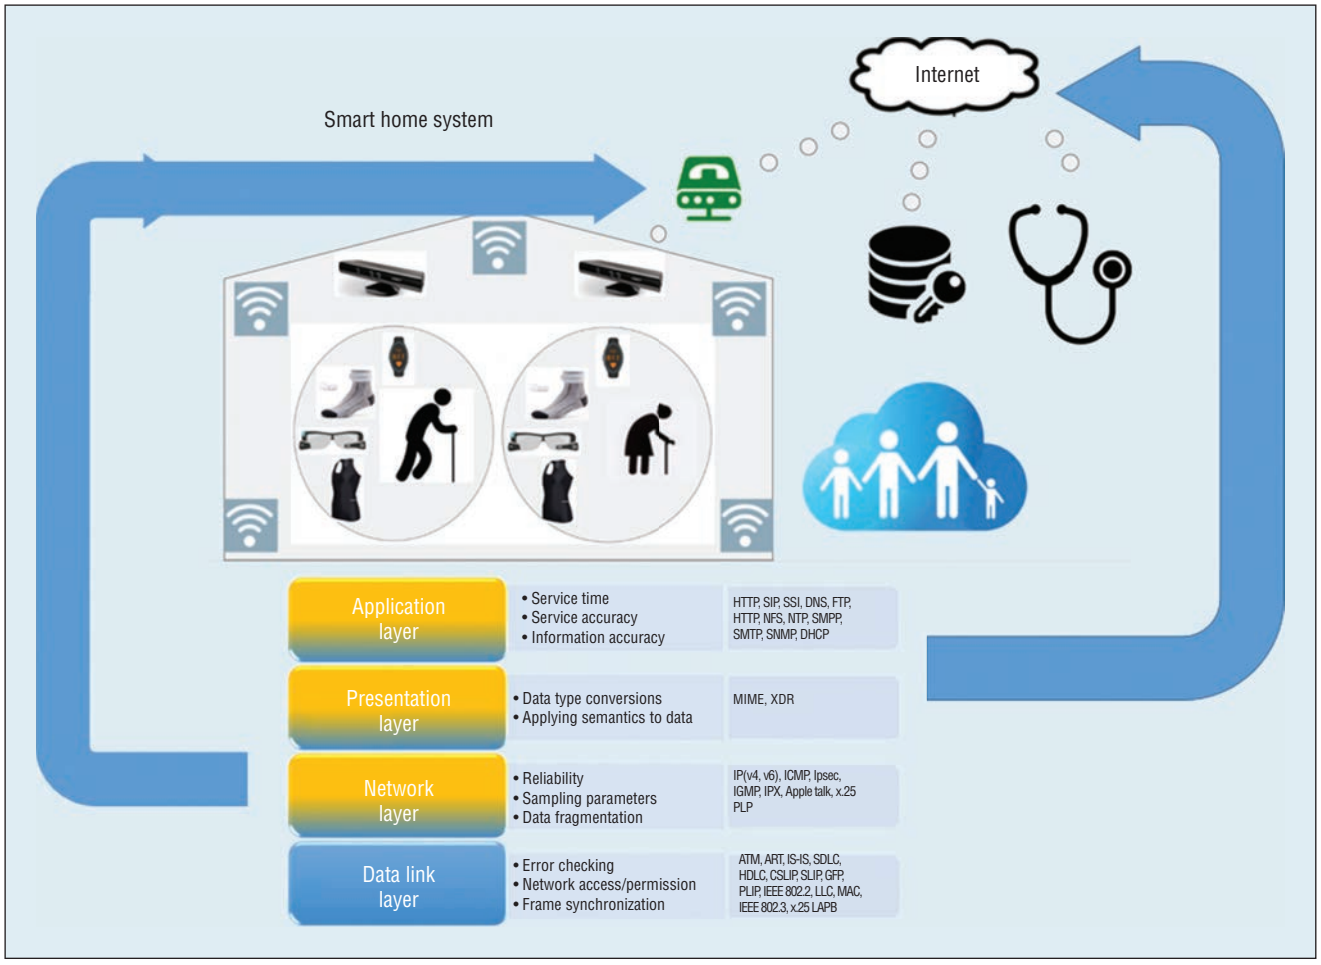


**Figure 3. A smart home system in which an older couple is monitored using wearables and environmental sensors for fall prevention and activity monitoring. The data are further transmitted via the Internet to enable real-time clinical decision support.**

unobtrusive, ambulatory health monitoring with instantaneous feedback to the user about the current health status. The devices are connected wirelessly via low-powered networking protocols such as Zigbee (motion sensors), Zwave, or Bluetooth (activity trackers), as well as through high-powered wired connections (Kinect). Device interoperability is crucial to ensure that all the data are recorded simultaneously, continually, and accurately. However, systems such as the semantic gateway bypass the network interoperability that acts as a bridge between the IoT devices and the Internet to allow part of the data processing to occur in the gateway, enabling faster decision support.[16]

## Effect of Application-Driven Quality of Service at the OSI Application Layer

A crucial and often overlooked challenge in terms of quality of service for the smart home system is that quality is dominated by its weakest link—that is, the lowest-quality sensor device. This could include a failed sensor, device-specific network connectivity issues, or even database malfunctions. As an example, consider the fall-detection system in the smart home setting. This system comprises heterogeneous sensors, such as depth and audio sensors, to detect falls inside the home and alert the clinicians. However, if the lowest frame rate among the sensing devices (say, the depth camera) is 2 frames per second, that will be the resolution of the overall system, conservatively speaking. Although this frame rate might be sufficient for activity-monitoring systems, the resolution could be too low if we want to detect falls occurring inside the home. Moreover, the depth sensor's low frame rate can seriously affect the fall-recognition system if it is a combination of multimodal sensors such as depth and audio, which relies on the sensor fusion for detecting the fall occurrence. To address this, we can use two factors for activity recognition: the individual sensors' data quality and knowledge of the activity itself.

An important point to note here is that both of the factors discussed in

Table 1. Factors affecting quality of service in smart home fall-detection systems.

| Factor | Information (example) | Solution |
|---|---|---|
| Activity understanding | A fall event corresponds to a "sudden change in height, with a sudden increased downward movement, as well as a corresponding trigger to the person being on the ground."[13] | Incorporating the semantic data quality check to look for the temporal sequence information of the activities. |
| Disparity in sensor data quality | The depth sensor is more prone to generate false alarms in detecting falls as compared to the motion sensors in the living area. | Weighted aggregation of the sensor devices depending on the location of the fall event. Fuzzy aggregation methods such as Sugeno and Choquet integrals[17] also allow uncertainty to be incorporated, which can take into account sensor data noise. |

Table 1 require prior knowledge of the activities, as well as evaluation of the sensor data quality that can be leveraged to learn the aggregation measures for multimodal data fusion.

## Internet of Everything or Indispensable Role of Humans in Quality Control

The two factors discussed in the quality of service of the smart home system in Table 1 are essential for the effectiveness of the smart home system. However, despite the incorporation of activity knowledge as well as the IoT device quality, the performance of the state-of-the-art fall-detection system is still low. One way to improve the system's performance is to incorporate human knowledge into the IoT system architecture. In fact, the Internet of Everything (IoE) is a concept that extends the IoT framework on machine-to-machine (M2M) communications to encompass people and processes for a much larger scale of data analytics. For our smart home-monitoring system, by incorporating a human-in-the-loop, the fall-detection system can achieve a much lower false-alarm rate that will alert the clinician and family members of a fall only after a technical nursing staff dedicated for this purpose has confirmed its occurrence. This can not only reduce clinician fatigue but also prevent undue panic through a more mediated IoE approach. Moreover, through an active learning process, the existing fall-detection system can update the algorithm to reduce the number of instances where the human-in-the-loop is required using the IoE design.

In a report from Cisco on IoE innovations,[18] a key insight was on the increased usage of mobile applications for interacting with IoE processes. For a fall-detection system, using a fall-detection mobile application will further allow clinicians and family members to access real-time feedback on their patient or loved one's status, thereby transforming the current clinical decision support system.

Overall, we see the effect of the IoT, and even the IoE, on two use cases in daily life. Through improved IoT data quality, the IoT can have a staggering impact on different facets of our existence, from entertainment, surveillance, transportation, and daily activities to industry applications and healthcare. ■


## Acknowledgments

We thank John Hughes from Wright State Internal Medicine and Mary Catherine Schafer from Jefferson City Medical Group for their valuable suggestions on the clinical aspects of the use cases described in this article.


## References


1. A. Sheth, "Internet of Things to Smart IoT Through Semantic, Cognitive, and Perceptual Computing," *IEEE Intelligent Systems*, vol. 31, no. 2, 2016, pp. 108–112.
2. J. Gubbi et al., "Internet of Things (IoT): A Vision, Architectural Elements, and Future Directions," *Future Generation Computer Systems*, vol. 29, 2013, pp. 1645–1660.
3. "ASQ Quality GlossaryQ," 2017; http://asq.org/glossary/q.html.
4. *ISO 9000:2005, Quality Management Systems—Fundamentals and Vocabulary*, ISO, 2005.
5. A. Sheth, P. Anantharam, and K. Thirunarayan, "kHealth: Proactive Personalized Actionable Information for Better Healthcare," *Workshop Personal Data Analytics in the Internet of Things*, 2014; http://knoesis.org/node/2237.
6. T. Banerjee et al., "Evaluating a Potential Commercial Tool for Healthcare Application for People with Dementia," *Proc. Int'l Conf. Health Informatics and Medical Systems*, 2015; http://corescholar.libraries.wright.edu/cgi/viewcontent.cgi?article=2442&context=knoesis.
7. L.J. Martin, "Aging Changes in Vital Signs," *Medline Plus*, 27 Oct. 2014; http://medlineplus.gov/ency/article/004019.htm.
8. R. Knudson et al., "Changes in the Normal Maximal Expiratory Flow-Volume Curve with Growth and Aging," *Am. Rev. Respiratory Disease*, vol. 127, no. 6, 1983, pp. 725–734.
9. A. de Pablo et al., "Pathophysiological Consequences of Lung Volume Reduction Surgery in Patients with Emphysema," *Archivos de Bronconeumología*, vol. 39, no. 10, 2003, pp. 464–468.
10. R. Pellegrino et al., "Interpretive Strategies for Lung Function," *European Respiratory J.*, vol. 26, no. 5, 2005, pp. 948–968.
11. M. Compton et al., "The SSN Ontology of the W3C Semantic Sensor Network Incubator Group," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 17, 2012, pp. 25–32.
12. T. Banerjee et al., "Recognizing Complex Instrumental Activities of Daily Living Using Scene Information and Fuzzy Logic," *Computer Vision and Image Understanding*, vol. 140, 2015, pp. 68–82.

13. T. Banerjee et al., "Monitoring Hospital Rooms for Safety Using Depth Images," *Proc. AAAI Fall Symp. Series AI for Gerontechnology*, 2012; www.eldertech.missouri.edu/wp-content/uploads/2016/07/Monitoring-Hospital-Rooms-for-Safety-Using-Depth-Images.pdf.
14. G. Alexander et al., "Density Map Visualization as a Tool to Monitor Activity Levels of Older Adults," *Gerontechnology*, vol. 9, no. 2, 2010, p. 186.
15. I. Taimiya Sylla, "Wireless Body Area Networks: What Engineers Need to Know," *EE Times*, 26 Sept. 2011; www.eetimes.com/document.asp?doc_id=1279106.
16. P. Desai, A. Sheth, and P. Anantharam, "Semantic Gateway as a Service Architecture for IOT Interoperability," *Proc. IEEE Int'l Conf. Mobile Services*, 2015, pp. 313–319.
17. L.M. Campos and M. Jorge, "Characterization and Comparison of Sugeno and Choquet Integrals," *Fuzzy Sets and Systems*, vol. 52, no. 1, 1992, pp. 61–67.
18. J. Bradley et al., *Internet of Everything in the Public Sector: Generating Value in an Era of Change*, Cisco, 2014.

**Tanvi Banerjee** is an assistant professor of computer science and engineering at Wright State University and Kno.e.sis. Her research interests include sensor validation and tying machine learning with sensor data for actionable information in healthcare applications, specifically in the domain of eldercare technologies. Contact her at tanvi@knoesis.org.

**Amit Sheth** is the LexisNexis Ohio Eminent Scholar, executive director of the Ohio Center of Excellence in Knowledge-Enabled Computing (Kno.e.sis) at Wright State University, and an IEEE Fellow. Contact him at amit@knoesis.org; http://knoesis.org.amit.

*This article originally appeared in* IEEE Intelligent Systems, *vol. 32, no. 2, 2017.*

Editor: Munindar P. Singh • m.singh@ieee.org

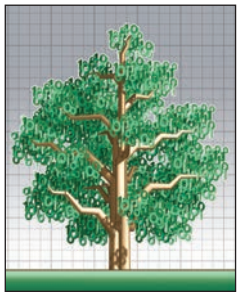

# Continuous Authentication and Authorization for the Internet of Things

**Muhammad Shahzad and Munindar P. Singh** • *North Carolina State University*

How can users be authenticated and authorized continuously for the Internet of Things, when most small smart devices lack the conventional interfaces used for authentication (such as keyboards, mice, and touchscreens)? Here, the authors explore potential solutions along with a related case study.

We're venturing into the era of the Internet of Things (IoT). As computing devices become smaller, smarter, and ubiquitous, computing has begun to embed into our environments by attaching to physical objects or things. IoT is bringing computing both onto our bodies and into our daily surroundings. Examples of on-body computing devices include human activity trackers, smart watches, and semi-permanent insulin pumps. Examples of in-environment computing devices include intelligent thermostats, smart appliances, remotely controllable household equipment, and weather-based automated lawn irrigation systems.

Although IoT devices often have compute power close to those of conventional computing devices from a few years ago, one of the ways in which typical IoT devices differ is that they lack conventional user interfaces in the form of keyboards, mice, and touchscreens. Examples of such computing devices include the Fitbit activity tracker, sewable computing devices such as the Arduino Lilypad, and smart fabrics. A motivation for eliminating such user interfaces isn't so much to reduce the cost as that the conventional interfaces often aren't appropriate for the intended applications. For example, a smart fabric can have embedded antennas and communicate information about the person wearing the fabric to devices such as smartphones, but it wouldn't quite make sense to attach a touchscreen to a shirt or a keypad to a Fitbit.

This lack of a user interface gives rise to a fundamentally challenging question: How do we authenticate and authorize users for the IoT, where we lack conventional user interfaces? For example, one of the latest features of the Apple Watch is that if a user owns a (sufficiently new version of) Macbook Pro, an iPhone, and an Apple Watch, the user can set up the Macbook Pro to automatically unlock without entering a password. More specifically, as soon as the user opens the lid of her Macbook Pro, the laptop automatically unlocks if the following four conditions are satisfied:

- The user is wearing the Apple Watch.
- The Apple Watch is connected to the user's iPhone via Bluetooth.
- The watch is in close proximity to the Macbook Pro.
- Either the iPhone or the watch has been unlocked at least once since the user last put on the watch.

This convenient feature carries a security threat, however. Suppose an attacker, possibly posing as a friend, gets hold of the user's watch and has physical access to her computer. Such a scenario might occur if the two are in a lunch meeting and the user steps away from the table to pick up something from the buffet, but leaves behind her watch and computer. If the attacker wears that watch and the user happens to unlock her phone while away from the table — but within the Bluetooth range of the watch — the watch unlocks as well. Then the attacker can use the watch to unlock the user's Macbook Pro without having to guess her password.

 Published by the IEEE Computer Society 

Although this technology employs an Apple Watch, which does have a conventional user interface in the form of a touchscreen, in principle, it can be extended to any wearable device with a Bluetooth interface, such as a Fitbit. These kinds of examples highlight the present challenge: How can we continuously authenticate a person using a device without a conventional interface? Here, we consider various solutions, including a case study for a Wi-Fi-based human authentication system (Wi-Fi uses radio frequencies near 2.4 and 5 GHz).

## Prospective Solutions

Because of the diversity of devices and applications, a universal solution to the problem of continuous authentication of users on devices without conventional interfaces might not exist. However, we can make progress by dividing IoT devices with which humans interact into two categories and studying solution directions for these categories separately. The first category consists of devices that maintain permanent physical contact with the user during usage, such as activity trackers, smart watches, and insulin pumps. The second category consists of devices that don't maintain permanent physical contact with humans, such as intelligent thermostats, occupancy sensors, and smart household appliances.

### Authentication on Devices That Maintain Continuous Physical Contact

Devices that maintain contact with the user can support new forms of biometric authentication. Most of these devices fall into two categories. Devices of the first category either contain an inertial measurement unit (IMU), which is comprised of an accelerometer and a gyroscope, or can have an IMU embedded quite easily. Devices of the second category contain a photoplethysmogram (PPG) sensor, which is comprised of a few (often two or three) LEDs and a few (again, often two or three) light sensors. Both IMUs and PPG sensors can enable user authentication.

Specifically, using the IMU, we can develop authentication techniques that are based on the principle that users frequently move their limbs in unique patterns throughout the time they use the device. An example of a well-known trait that differs across users is gait. If we can extract the patterns in the output of an IMU sensor due to the user's unique gait, we can use simple machine learning techniques to learn these patterns and apply them to continuously authenticate the user based on gait. If such a technique was developed and put into practice, a device of the first category (worn by the user) could monitor the user continuously and frequently authenticate the user's legitimacy before allowing the user to perform appropriate operations. For example, the watch in the Macbook Pro setting wouldn't authenticate the attacker, which would prevent the Macbook Pro from being spuriously unlocked. Several behavioral biometrics solutions have been proposed that employ the IMU to authenticate users.[1,2]

Similarly, the PPG sensor provides an opportunity to study the PPG signal for unique patterns in blood flow rhythm. Researchers have shown that due to slight variations in every human's heartbeat rhythm, echocardiogram (ECG) signals contain small information — but this is enough to indicate what's unique to an individual user.[3,4] Consequently, just by using the ECG signal, we can design user schemes to authenticate users. Although several ECG-based user authentication systems have been proposed, this technology has yet to achieve sufficient effectiveness to see widespread deployment. Because the PPG signal is generated based on the amount of blood flow in the user's veins, which depends on how the user's heart pumps blood, the PPG signal could contain enough information to enable user authentication. Using the PPG signal is particularly challenging, however, because this signal is sensitive to the motion of a person's limbs: that is, the PPG measurement depends on the person's speed of movement. Fortunately, most devices these days that come with a PPG sensor also come with an IMU. Therefore, we can potentially use information from the IMU sensors to measure the amount of motion of the limb and combine the two signals — or correct the measurement of the PPG sensor — to enable authentication.

### Authentication on Devices That Don't Maintain Permanent Physical Contact

A more challenging problem is to design an authentication scheme that can identify users for devices that don't maintain permanent contact with users. Such devices include those embedded into our environment. For example, consider an application that integrates a user's calendar with the user's home lighting and is controlled with speech. Whenever a calendar generates a notification for the user, the user's location is automatically determined through proximity or movement sensors, and lights of the appropriate room are flashed to alert the user. Suppose this application is enabled or disabled through voice commands from a specific user. Then, an attacker could replay previously recorded voice commands of the original user. That is, voice-based authentication is insufficient. We need effective methods to continuously and unintrusively authenticate users without, for example, requiring the user to wear sensors such as IMUs.

A potential approach for developing such an authentication system is to employ pervasive modalities such as radio frequency (RF) signals, ambient light, and sound, which are present all around us. The intuition behind RF-based authentication is

that the wireless channel metrics — such as channel state information (CSI) and received signal strength (RSS) — change based on a user's presence and movement. The patterns of change in these metrics depend on the way the user moves. Because different users have different gaits, they produce different patterns of change in wireless channel metrics. An RF-based user authentication system can apply machine learning techniques to associate each user with his or her patterns of change and identify the user at runtime based on the learned associations.

Specifically, with a human walking around, because a human is mostly made of water, the Wi-Fi signal reflected by the human body generates unique, although small, variations in CSI measurements on the receiver due to the well-known multipath effect of wireless signals. These variations in CSI enable signal processing techniques to obtain gait information such as walking speed, gait cycle time, footstep length, and movement speeds of legs and torso. Because each human has a unique gait, the gait patterns that the Wi-Fi receiver obtains can be used to recognize a walking human subject.

Similarly, the intuition behind light-based authentication is that as a user moves in an indoor environment, the amount of light he or she reflects and blocks depends on his or her patterns of movement. As different users have different gaits, the patterns of change in intensity of light, as measured by light sensors deployed on the floor, are also different. A light-based user authentication system can learn these patterns and apply them to identify users at runtime. A similar intuition holds for audio-based user authentication.

## Authorization in the IoT

So far we've talked of authentication because it provides concrete use cases. But authentication by itself is usually meaningless. The point of authentication is to provide a basis for making a decision — about which resources to provide access to which person for what purpose and when. In broad terms, we aren't so much interested in seeing who specifically is around but what information or device to share with that person under what circumstances.

Consider a situation where a user is wearing various health-monitoring devices, including an ECG reader. Ordinarily, the data gathered by such devices would be confidential. Now, suppose the user is having a serious medical problem, such as a heart attack. In such a case, it might be acceptable behavior for the relevant software application to reveal the data from his ECG reader as well as data about recent physical activity to anyone who is present nearby and might be willing and able to help. But if the user is already in a hospital, then perhaps the application doesn't need to be quite so forthcoming in revealing its user's information to strangers. That is, here the decision changes from a focus on authentication of the counterparty to determining some attributes of the information resource and of the current context. Indeed, modern approaches[5] to authorization express policies in terms of attributes of principals, resources, and contexts instead of specific identifiers or roles. The IoT can readily accommodate such approaches by accumulating a rich variety of attribute values from the available devices.

In this example, the decision about whether to share some data is based upon the data values themselves (for example, sharing with anyone if the ECG indicates distress), but in general the decision might be based on the totality of available information. In particular, we can have situations where an application grants access to wearable devices based on environmental devices or the other way around. For example, in a home eldercare setting, if the environmental sensors (whether Wi-Fi or light-based) indicate a lack of movement for a prolonged period, the eldercare application might disclose data from wearable devices that capture the resident elder's health condition. Conversely, the data from a wearable sensor being anomalous might lead the application to verify whether a qualified caregiver was currently in the same room as the resident.

## Case Study

To validate the effectiveness of such an approach to leverage variations in pervasive modalities, working with colleagues, we developed a Wi-Fi-based human authentication system, called WifiU, which recognizes users based on their gait.[6] We developed WifiU entirely using COTS Wi-Fi devices to capture fine-grained gait patterns. WiFiU consists of two Wi-Fi devices: one for continuously sending signals, which can be a router, and one for continuously receiving signals, which can be a laptop. In WifiU, the receiver measures channel state information (CSI) of each received Wi-Fi frame. Fundamentally, WifiU recognizes humans based on who they are, because WifiU extracts unique biometrics information from Wi-Fi signals and uses it to perform human authentication.

Compared with traditional gait-recognition systems, which use cameras, floor sensors, or wearable sensors to capture gait information, WifiU is easier to deploy and has better coverage. From the deployment perspective, WifiU doesn't require any special hardware (such as floor sensors) and doesn't require the human subject to wear any hardware (such as an IMU). Wi-Fi devices are ubiquitous and most homes and offices are covered by Wi-Fi signals. The hardware that we experimented with — namely a NetGear JR6100 Wi-Fi router and ThinkPad X200 laptop (with an Intel 5300 WiFi NIC) — required no modifications.

Furthermore, unlike cameras, WifiU doesn't require lighting and works in the dark just as well as in bright light.

In designing WifiU, we faced many technical challenges. For example, it's nontrivial to profile gait patterns using CSI dynamics. Extracting gait information from CSI signals is difficult, because the signal reflections of different body parts are mixed together in the CSI waveform. As different human body parts move at different speeds while walking, the radio signal reflections from different body parts have different frequencies. To separate the radio signal reflections from the different body parts, we convert CSI waveforms (of two dimensions: time and amplitude) into spectrograms in the time-frequency domain (of three dimensions: time, frequency, and amplitude). We apply spectrogram enhancement techniques to reduce signal noise. The resulting spectrograms yield detailed human gait information similar to those generated by Doppler radars.

We conducted experiments on WiFiU using our gait database that contains more than 2,800 gait instances collected from 50 human subjects walking in a typical laboratory with an area of 50 square meters (see Figure 1). We anonymized all collected data to protect participants' privacy. Over the 50 subjects, WifiU achieves recognition accuracies of 79.3, 89.5, and 93.0 percent for the Top-1, Top-2, and Top-3 candidates, respectively. (Here, Top-*N* means that one of the selected *N* persons is the person who truly generated that gait observation.)

With the current implementation using a single wireless link, WifiU has three limitations. First, the distance between the human subject and the Wi-Fi devices is limited to six meters. To address this limitation in future work, we can deploy multiple Wi-Fi sender-receiver pairs in the target area. Second, the recognition accuracy is limited to 92.3 percent for Top-1 candidates. Whereas this accuracy might be acceptable for many personalized services such as adjusting room temperature and background music in smart buildings, it might not be high enough for settings that require high accuracy, such as accessing your email. Third, the number of walking human subjects is limited to one. In practice, the targets should walk along a given path one by one to ensure good recognition performance, as is the case for an airport security check. To address this limitation in future work, we plan to use multiple Wi-Fi receivers to separate signals of multiple humans using the differences in received signals at multiple receivers.

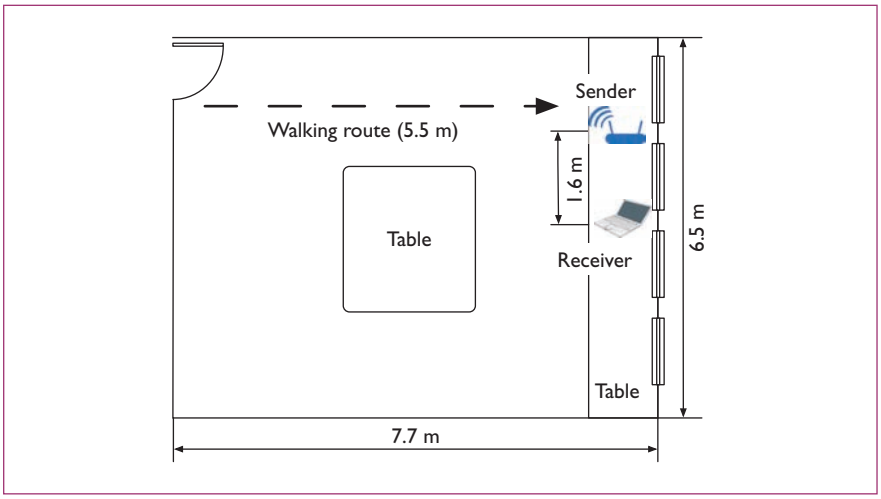


*Figure 1. Data collection environment. Walking in an area of 50 square meters, we gathered more than 2,800 gait instances from 50 human subjects.*

The IoT is in a nascent stage, but the arc of technology and the potential benefits it offers suggest that the IoT's presence will only increase. By placing people in information-rich environments, especially those that are natural and feel natural, the IoT exposes users to new security and privacy threats. It simultaneously demands stronger (that is, continuous) authentication and authorization and takes away conventional information modalities. Fortunately, the IoT provides new ways to address these challenges through innovative uses of technology. Therefore, the prospects of unintrusive authentication and authorization leading to context-sensitive policies are encouraging.

However, these unintrusive authentication technologies create potential privacy threats through the infrastructure in that an attacker who can obtain access to the infrastructure might apply these techniques without the user being aware of having been identified. For example, an attacker can potentially read Wi-Fi signals to identify victims without being detected. Consider a scenario where a burglar attempts to figure out who is at home by eavesdropping on the Wi-Fi signal emitted by the Wi-Fi router in the victim's house. As Wi-Fi signals can penetrate through obstacles such as furniture, wooden doors, and walls, the burglar needs to only passively measure the CSI of the signal outside the house without needing to decode the Wi-Fi packets' content. Therefore, it would be difficult for the victim to prevent certain breaches of privacy. Although avoiding privacy breach isn't the focus of this article, we hope this work highlights this privacy risk to the research community and encourages future work. A previous column[7] addresses the privacy risks in intelligent user interfaces, some of which might be exacerbated in combination with the IoT.

**Acknowledgment**

Munindar Singh thanks the US Department of Defense for support through the Science of Security Lablet.

**References**

1. R. Mayrhofer and H. Gellersen, "Shake Well before Use: Authentication Based on Accelerometer Data," *Proc. Int'l Conf. Pervasive Computing*, 2007, pp. 144–161.
2. T.T. Ngo et al., "The Largest Inertial Sensor-Based Gait Database and Performance Evaluation of Gait-Based Personal Authentication," *Pattern Recognition*, vol. 47, no. 1, 2014, pp. 228–237.
3. S.I. Safie, J.J. Soraghan, and L. Petropoulakis, "Electrocardiogram (ECG) Biometric Authentication Using Pulse Active Ratio (PAR)," *IEEE Trans. Information Forensics and Security*, vol. 6, no. 4, 2011, pp. 1315–1322.
4. Z. Zhang et al., "ECG-Cryptography and Authentication in Body Area Networks," *IEEE Trans. Information Technology in Biomedicine*, vol. 16, no. 6, 2012, pp. 1070–1078.
5. D.R. Kuhn, E.J. Coyne, and T.R. Weil, "Adding Attributes to Role-Based Access Control," *Computer*, vol. 43, no. 6, 2010, pp. 79–81.
6. W. Wang, A.X. Liu, and M. Shahzad, "Gait Recognition Using WiFi Signals," *Proc. Int'l Conf. Pervasive and Ubiquitous Computing*, 2016, pp. 363–373.
7. C.J. Hazard and M.P. Singh, "Privacy Risks in Intelligent User Interfaces," *IEEE Internet Computing*, vol. 20, no.6, 2016, pp. 57–61.

**Muhammad Shahzad** is an assistant professor in the Department of Computer Science at North Carolina State University. His research interests include the measurement, networking, and user interface aspects of the IoT as well as measurements, design, and modeling of computer networks. Shahzad has a PhD in computer science from Michigan State University. Contact him at mshahza@ncsu.edu.

**Munindar P. Singh** is a computer science professor at North Carolina State University. His research interests include the conception, engineering, and governance of sociotechnical systems as a way to tackle concerns such as security and privacy. Singh is a Fellow of IEEE and the American Association for Artificial Intelligence (AAAI), a former Editor in Chief of *IEEE Internet Computing*, and the current Editor in Chief of *ACM Transactions on Internet Technology*. Contact him at singh@ncsu.edu.

*This article originally appeared in* IEEE Internet Computing, *vol. 21, no. 2, 2017.*

# Visual IoT: Architectural Challenges and Opportunities

**Ravi Iyer**
Intel

The emergence of ultra-low-power sensing devices along with connectivity to gateways and cloud services has led to an end-to-end Internet of Things (IoT) architecture for many real-world usages. Visual IoT is one such class of IoT that poses significant end-to-end challenges due to the need for sensing and processing of visual data. The richness of visual data provides many opportunities for analytics, while at the same time requiring high computational capabilities and therefore potentially high-bandwidth data transfer to a more powerful node in the end-to-end architecture. Memory and storage needs are also more pronounced in visual IoT solutions, requiring careful thought to developing an intelligent memory hierarchy for visual storage and retrieval. In this article, I examine the computing, memory, and interface implications for end-to-end visual IoT architectures and discuss potential solutions and tradeoffs in each of these areas. Before we start, let's go over a brief overview of visual IoT usage domains.

## Visual IoT Overview

Beyond photography, cameras have been used widely in multiple domains, ranging from security (for example, surveillance and monitoring), entertainment (recording of public and personal events, such as sports and music), and, more recently, interactive environments (augmented, virtual, and merged reality; see www.intel.com/content/www/us/en/architecture-and-technology/virtual-reality-overview.html) and robotics and drones[1] (navigation, delivery, interaction, and assistance). With the emergence of depth-sensing cameras, such as Intel RealSense (www.intel.com/content/www/us/en/architecture-and-technology/real-sense-overview.html), analysis of the captured visual scene becomes even more attractive for many of these scenarios.

In many of these scenarios, three types of platforms compose the end-to-end IoT architecture (see Figure 1):

- visual sensing nodes that capture the data and potentially do some local processing;
- gateways, phones, or on-premise platforms that can stage the data and provide higher computing capability; and
- cloud servers that provide services for search, analytics, or simply storage.

Much like real estate, the key to efficiently architecting a visual IoT architecture is location (where to perform the computation), location (where to store the data), and location (where to enable interfaces and tools for analytics). Let's start by examining the computing location challenge and then move to memory and interfaces.

**Editors' note:** We invited two industry experts to discuss the opportunities and challenges surrounding the Internet of Things. The following are their views on the topic. —*Vijay Janapa Reddi and Hyesoon Kim*

## Computing in Visual IoT: Partitioning and Heterogeneity

Partitioning the work in an end-to-end visual IoT architecture is a challenge, because it requires the balancing of multiple important dimensions:

- the sensing node's battery life,
- the latency of the interaction,
- throughput benefits on the server versus bandwidth costs of the transfer, and
- security and privacy implications of the data.

The partitioning problem is essentially across heterogeneous platforms as well as within heterogeneous processing elements (cores versus GPUs versus accelerators) within a platform. Also, the application can dictate the partitioning strategy in some cases, in which a subset of operations on the sensor node can deliver some minimal useful experience, while the processing at the server is used for heavier computations.

Let's take an example scenario of a visual agent monitoring a home environment. The key aspects of a home agent include

1. anomaly detection,
2. saliency and summarization,

2469-7087/17/$33.00 © 2017 IEEE Published by the IEEE Computer Society

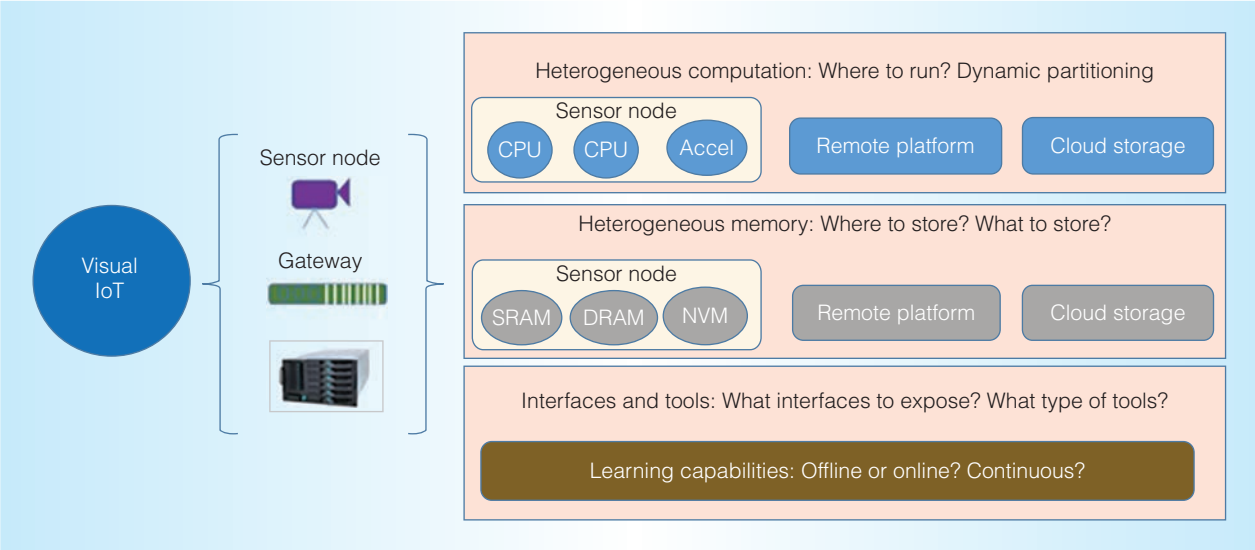


Figure 1. Visual IoT. Example implications on architecture research include how to dynamically partition the computation across the heterogeneous architectures, manage the memory across the end-to-end system, and integrate offline/online learning capabilities and tools.

3. detection of patterns of behavior, and
4. recognition and interaction with a person in the home from a Q&A standpoint.

The scenario becomes even more complicated if the visual agent is mobile (like a robot) versus a static visual agent in which the backgrounds can be predetermined. For aspects 1 and 4, the response time is critical, so local processing is desirable, whereas for aspects 2 and 3, batch processing is more useful because of the large amount of data needed before processing.

Such scenarios are common and require careful examination of whether the processing can be performed on the computing core of the sensor node itself, offloaded to a local accelerator, sent to a gateway within the home, or offloaded to a private cloud where the analysis can be accomplished. A static solution would end up determining how to employ the most efficient engine (fixed function or configurable accelerator) at each node in the end-to-end architecture. Instead of statically determining this heterogeneous architecture and partitioning balance, a dynamic partitioning solution is even more suitable if the solution has to be customized for different homes and similar environments. As a result, solutions such as remote offloading are becoming more important from a flexibility and customization point of view.[2] Research and development in heterogeneous architectures with partitioning capabilities that retain flexibility while maximizing efficiency and customizability will continue to be predominant for visual IoT, as well as other rich environments.

## Memory in Visual IoT: Saliency, Storage, and Hierarchy

Another major challenge in the end-to-end visual IoT architecture is management of the visual data. Although the richness of visual data is attractive, it is also true that much of the visual data captured can be potentially discarded, and only a summary typically needs to be retained. The key is to figure out what visual data summary must be retained by potentially extracting the salient segments of the visual stream. The basis for saliency depends entirely on the usage in question. For example, in the home scenario, the salient aspects might be the key activities that happened throughout the day and the anomalies and novel occurrences that were identified. In a recent study,[3] the authors demonstrated the ability to summarize a video by optimizing for similarity and coverage. Analyzing such algorithms and capabilities and converting them into appropriate computing and memory implementations[4] is a critical area of research for future visual sensor nodes, as well as the gateways and servers that maintain the data.

Beyond saliency and summarization of frames, it is also critical to identify key entities and activities in visual data to enable fast search and indexing. The question becomes what metadata needs to be extracted and where such metadata should be stored (on the sensor node, in the gateway, or on the cloud server). In addition, there is a question of what type of memory is most suitable for the metadata in question. This calls for an end-to-end heterogeneous memory architecture consisting of different memory types, ranging from cache to DRAM to nonvolatile memory to storage. Identifying the right balance of such heterogeneous memory across each of the nodes in the end-to-end architecture is critical as visual IoT usages explode and cause bandwidth challenges for retrieval of data.

### Interfaces and Tools for Visual IoT: Learning and Development

Finally, it is important to consider appropriate interfaces for visual IoT platforms. For example, as machine learning techniques get adopted to analyze sensor data, it becomes important to understand how to take advantage of both offline and online learning techniques. As an example, if the visual agent wants to understand gestures made by the people in a home, it is extremely useful to enable interfaces and tools that allow the agent to train on and download these capabilities. By making such capabilities broadly available, developers will be able to provide many analytics capabilities employing rich sensor data and potentially crowdsourced training data. Especially as sensor nodes become more capable (such as the Intel Curie Module[5] with pattern matching capability), new tools that enable developers to use such capabilities (such as the Intel Knowledge Builder toolkit [http://software.intel.com/en-us/intel-knowledge-builder-toolkit]) are critical for the rapid deployment of IoT solutions.

Visual IoT is a rapidly growing class of usages with the proliferation of smart cameras with increasing capabilities. Future areas of research include developing heterogeneous architectures and dynamic partitioning capabilities across end-to-end visual IoT, examining heterogeneous memory stores for visual data management and retrieval, and tools and interfaces for fast deployment of analyzing visual and other IoT solutions. MICRO

### References


1. K. Kaplan, "The Future of Drones: Market Prepares for Takeoff," Intel, Sept. 2016; http://iq.intel.com/drone-economy-prepares-takeoff.
2. H. Eom et al., "OpenCL-Based Remote Offloading Framework for Trusted Mobile Cloud Computing," *Proc. Int'l Conf. Parallel and Distributed Systems*, 2013, pp. 240–248.
3. S. Chakraborty, O. Tickoo, and R. Iyer, "Adaptive Keyframe Selection for Video Summarization," *Proc. IEEE Winter Conf. Applications of Computer Vision*, 2015, pp. 702–709.
4. T. Lee et al., "Low-Complexity HOG for Efficient Video Saliency," *IEEE Int'l Conf. Image Processing*, 2015; doi:10.1109/ICIP.2015.7351505.
5. "Intel Curie Module & Intel IQ SW Fact Sheet," Intel, Aug. 2015; www.intel.com/content/www/us/en/wearables/intel-curie-fact-sheet.html.



**Ravi Iyer** is a senior principal engineer, CTO, and director in New Business Initiatives at Intel. He is an IEEE Fellow. Contact him at ravishankar.iyer@intel.com.


# Toward a Self-Learning and Energy-Neutral IoT


**Emre Ozer**
ARM


A typical Internet of Things (IoT) device comprises five components: sensor, microcontroller, memory, battery/energy harvester, and radio. It is a device that collects, preprocesses, stores, and transfers data received from a sensor to a host (for example, a reader via RFID, a smartphone via Bluetooth, or the cloud through a gateway) wherein data processing is performed. The microcontroller is mainly responsible for control and simple data preprocessing, and the radio is used to transmit short data packets. Hence, the battery can last for years before it is recharged or replaced.

Such a simple IoT device is no longer adequate because emerging applications (such as medical, structural/environmental monitoring, and e-textiles) demand ambient intelligence or cognition and real-time response from IoT devices. A new class of IoT devices called *self-learning IoT devices* will emerge to provide cognitive services, such as situational awareness, anomaly detection, activity, and pattern and emotion recognition, which are essentially machine learning algorithms. Real-time response from self-learning IoT devices is needed because an anomaly or critical activity must be detected or recognized in situ and reported immediately, because transmitting the sensor data via radio to its host to do this will be costly in terms of energy and latency.[1] For this reason, the cognitive action must take place in the device, not in the host. For example, an implantable chip must detect an abnormal condition in the organ and must take an action in real time. It cannot afford to wait for a critical decision to be made by the host.

A self-learning IoT device will accommodate multiple sensors and a more powerful computation engine to perform

computationally intensive sensor fusion and to run machine learning algorithms. It will need a good size on-device memory (SRAM and nonvolatile memory) to store the code and buffer the streaming sensor data before and after data fusion. Integrating multiple sensors, a relatively higher-performance computation engine, and more on-chip storage in a self-learning IoT device will consume more energy than a simple IoT device, and will put incredible pressure on the battery. Self-learning IoT devices will be deployed in such environments in which recharging or replacing the battery is not possible—for example, an IoT device implanted in a human body, integrated into a building's foundation, or embedded in the textile fabric. Hence, the battery in the device must operate for a long time (for example, more than 10 years) and be charged by multiple energy harvesters that are integrated into the device to harvest ambient energy (such as thermal, vibration, solar, pressure) in order to charge the battery. This is a concept called *energy neutrality,*[2] in which the battery will always be charged by energy harvesters in the device[3] and should never be recharged by human intervention.

The next phase in the IoT's evolution is self-learning and energy-neutral devices having the properties of cognition, real-time response, and perpetual energy. The main challenge is to run computation and memory-intensive sensor fusion and machine learning algorithms in a device powered only by the harvested energy. This opens up opportunities to design novel computation engines, memory subsystems, and energy management units, considering not only energy efficiency but also energy neutrality. The computation engine in the device must be equipped with single-instruction, multiple data/digital signal processing capabilities and be coupled with one or more machine learning hardware accelerators (such as a deep neural network). Alternatively, the computation engine can be tightly coupled with sensors that will stream data directly to the computation engine, such that it may have an analog preprocessing front end tightly coupled to the sensors, and the machine learning algorithms will run on the engine's digital back end engine. Today, state-of-the-art microcontrollers have up to 4 Mbytes of flash memory and much less static RAM, and future IoT devices will not have orders of magnitude larger on-chip storage because of the cost issues. Machine learning algorithms—in particular, deep neural networks—take up a significant storage space, so it will be a challenge to store a large number of network parameters on chip. Hardware and software compression techniques will be used to deal with the large parameter space in deep neural networks. Also, alternative machine learning algorithms that are more adaptive, resource efficient, and energy efficient (such as nonparametric Bayesian methods[4]) can be developed for self-learning and energy-neutral IoT devices. The management of the harvested energy is a critical process, and the data must be sensed, fused, stored, and processed, and the response given, before the harvested energy in the battery depletes. The harvested energy management must be performed by a combination of innovative software and hardware techniques, such as the prediction of the harvested energy before task execution, or new instructions to control the energy harvesting process.

Self-learning and energy-neutral IoT devices will also emerge in the printed electronics world.[5] Printed electronics offers cost-effective fabrication of electronics with low-cost substrates and materials (such as plastic and paper), simpler processing and patterning steps, and disposability. It has found applications in sensors, RFIDs, solar cells, batteries, and displays in the fields of medical, wearable, textile, automotive, and packaging applications. Smart printed devices have already been demonstrated as smart tags, labels, packages, e-textiles, and wearables. For example, T.E. Halterman built a printed alarm armband that monitors the vital signs of hospital patients.[6] The flexible armband contains a solar panel, piezoelectric speaker, temperature sensor, and power supply circuit, all of which are organic components in a wearable form factor. It is self-powered by the solar panel, and the speaker sounds an alarm when the temperature sensor measures a temperature between 36.5 to 38.5 degrees Celsius. These early demonstrators are the precursors of future self-learning and energy-neutral printed IoT devices. The main advantage of printed electronics is that they allow low-cost customization thanks to the low-cost flexible substrate and materials, and do not require costly clean rooms, unlike silicon. This offers a unique opportunity, in particular, to customize the computation engine to the needs of the cognitive application that will be running on the device. For example, an energy-efficient support vector machine (SVM) can be designed as a custom computation engine (rather than using a less energy-efficient general-purpose computation engine) and printed for a single-use smart packaging product, because it will run only the SVM. This will not be possible in silicon, because customization (that is, ASIC) is extremely costly. Thus, printed electronics will pave the way to low-cost customization of efficient computation engines for future printed self-learning and energy-neutral IoT devices.

Future IoT devices will become more intelligent and aware of their environment, and will integrate more capable computation engines to perform cognitive activities. However, these devices will still be constrained by energy efficiency and limited energy capacity, as in today's dumb IoT devices. They will be so deeply embedded that they will not be accessible to replace or recharge their batteries, and will have to depend on energy harvesters to become self-sustained or energy-neutral. This will be even more prominent for printed electronic devices that will be manufactured for a single use. The main engineering

challenge is to design an energy-neutral device that will be deployed in critical missions and stay operational for a long time, but at the same time run computationally complex machine learning algorithms. Nevertheless, this challenge brings up unique opportunities for system architects, designers, and software developers to come up with holistic solutions not only for the self-learning and energy-neutral IoT devices in silicon, but also in emerging printed electronics. MICRO

## References


1. R.C. Carrano et al., "Survey and Taxonomy of Duty Cycling Mechanisms in Wireless Sensor Networks," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 1, 2014, pp. 181–194.
2. M. Magno et al., "Infinitime: A Multi-sensor Energy Neutral Wearable Bracelet," International Green Computing Conference (IGCC), 2014.
3. A.S. Weddell et al., "A Survey of Multi-source Energy Harvesting Systems," *Proc. Conf. Design, Automation and Test in Europe*, 2013, pp. 905–908.
4. Y. Raykov et al., "Predicting Room Occupancy with a Single Passive Infrared (PIR) Sensor through Behavior Extraction," *Proc. ACM Int'l Jt. Conf. Pervasive and Ubiquitous Computing*, 2016, pp. 1016–1027.
5. S. Khan, L. Lorenzelli, and R. Dahiya, "Technologies for Printing Sensors and Electronics over Large Flexible Substrates: A Review," *IEEE Sensors J.*, vol. 15, 2015, pp. 3164–3181.
6. T.E. Halterman, "Flexible, 3D Printed, Solar Powered Thermal Alarm for Patient Monitoring," *3D Print*, 26 Feb. 2015; http://3dprint.com/47116/3d-printed-fever-alarm.



**Emre Ozer** is a principal research engineer at ARM Research in Cambridge. Contact him at emre.ozer@arm.com.



*This article originally appeared in* IEEE Micro, *vol. 36, no. 6, 2016.*

# Osmotic Flow: Osmotic Computing + IoT Workflow

**Matteo Nardelli**
University of Rome Tor Vergata

**Stefan Nastic and Schahram Dustdar**
TU Wien

**Massimo Villari**
University of Messina

**Rajiv Ranjan**
Newcastle University

The rapid evolution of Internet of Things (IoT) devices (e.g., sensors and gateways) and the almost ubiquitous connectivity (e.g., 4G, Wi-Fi, RFID/NFC, Bluetooth, IEEE 802.15.4) are forcing us to radically rethink how to effectively deal with massive volume, velocity, and variety of big data produced by such IoT devices. There are currently 6.4 billion IoT devices in use around the world and their number, capabilities, as well as the scope of their use, keeps growing rapidly. According to Gartner (http://www.gartner.com/newsroom/id/3165317) the number of IoT devices will reach 20.8 billion by 2020, and, by then, IoT service spending will reach $1,534 billion and hardware spending $1,477 billion.

As IoT expands into various application domains such as healthcare, utility grids, cities, agriculture, transportation, industry 4.0, and disaster management, need for investigating on-the-fly computation over the IoT data streams is ever more pressing. Indeed, most IoT applications are modeled as data transformation workflows that consists of: i) multiple interdependent, heterogeneous data analysis computational and programming models that realise various data transformation tasks from data ingestion to analysis, ii) virtualised/non-virtualised computational and network infrastructure, iii) communication media of various kinds (including wireless). Currently, powerful Cloud Datacentres (CDCs, e.g. AWS[1]) provide computation and data storage resources for IoT workflows, but they suffer from limited bandwidth and network latency, and support neither latency-sensitive applications nor applications that rely heavily on the data streaming from IoT data sources for computing intelligence in real-time (in the form of data ingestion and data analysis).

A possible solution to augment the scalability of CDCs lies in taking advantage of the ever-increasing computational and storage capabilities available at the network edge.[2,3,4] We note in the previous instalment of "Blue Skies" sensing and networking devices available at the network edge constitute a new type of computing infrastructure, the Edge Datacentre (EDC).[5] An EDC may vary in scope and capability, including gateways (Raspberry Pi 3, UDOO board, esp8266, Meshlium Xtreme, Arduino), software defined network solutions (e.g. Cisco IOx), or smart phones equipped with sensors. To facilitate highly distributed and federated computing environments, we proposed Osmotic Computing paradigm[5] that enables the automatic deployment of microservices over inter-connected EDC and CDC. The benefits of integrating EDC and CDC has already been recognised by several companies and open source initiatives, in-

 Published by the IEEE Computer Society 

cluding CISCO, AWS[1], and Google[3], and the OpenFog Consortium.[4] For example, AWS has enriched its CDC offerings with near-edge computing and storage capabilities (i.e., Snowball Edge, Greengrass).

Nevertheless, the usage of an Osmotic Computing infrastructure (CDC+EDC) poses new challenges for IoT workflow application developers and operations managers as they need the awareness of resource/device (CDC server vs. IoT gateway) heterogeneity, virtualisation software heterogeneity (e.g., hypervisor vs. container), data analytic programming model heterogeneity (stream processing vs. batch processing), geographic distribution, and network performance uncertainties.

Existing streaming data analysis platforms including (e.g., Spark[6], Heron[7], Google Dataflow[8], AWS, Kinesis[1], StreamCloud, Apache Storm), are CDC-centric, hence they do not meet the resource management and scheduling requirements for IoT workflows that require coordinated mapping for data analysis activities to both CDC and EDC. Many workflow application management platforms such as Pegasus, Triana, Taverna, Galaxy, e-Science Central, and Kepler support the development, deployment and execution of scientific workflow applications on CDC without considering newly evolved EDC capabilities. Apache Oozie and Linkedin Azkaban support a Hadoop workflow, but in a rather rigid manner that works well for only batch processing activities. Data analytics platforms such as YARN, Mesos  Amazon IoT and Google Cloud Dataflow can support manual provisioning of multiple data transformation tasks on CDC resources, but only in a performance-agnostic way.

## The Osmotic Flow Model

We propose *Osmotic Flow*, a new model for holistically programming, mapping and executing IoT data transformation workflow applications on a distributed infrastructure combining both EDC and CDC resources. In the Osmotic Flow model, an IoT workflow application is modelled as a directed (potentially cyclic) graph with data transformation tasks as its nodes, and dataflow dependencies (or control flow dependencies for computational synchronization, if/as needed) between data transformation tasks as its vertices. Osmotic Flow model permits data transformation tasks to be distributed, managed, and executed across any available CDC and EDC provider. A data transformation task encapsulates a microservice (e.g., Docker, Unikernel), a computational model (e.g. statistics, clustering, classification, anomaly detection, accumulation), and a data analysis programming model (e.g., stream processing, batch processing, SQL, NoSQL, data ingestion).

### Motivation

Let us consider a contemporary Smart City, where a plethora of IoT sensing devices with Internet connectivity are disseminated all over the urban environment. Buses, trains, and taxis continuously communicate their position; vehicles notify congested routes; citizens often geo-locate their position in messages, photos, videos, or accessing specific services.

All these IoT data sources continuously produce ever-increasing streams of data that can be collected and processed to get the so-called "pulse of the city", thus fostering awareness and capacity of taking informed decisions. Intelligent services aimed at improving the citizens' quality of life can be built by merging, filtering, correlating, and transforming these diverse data streams.

For example, a smart traffic light IoT application (see Figure 1) can identify traffic congestions and proactively change traffic light priorities and speed limits, so to reduce ripple effects and relieve the environmental impact. Let us focus on a single road divided in multiple segments and managed by traffic light. The traffic light is instrumented with appropriate IoT sensor (e.g., light state sensor, CCTV) and actuator. The traffic light sensor produce data streams about their current state (i.e., color of light turned on and light change timings) while the CCTV (traffic congestion indicator) produces visual evidence of congestion. The smart traffic light IoT application assumes that traffic congestion is directly proportional to the number of cars queued at each intersection and inversely proportional to the average speed per road segment, each road segment is equipped with above IoT sensor and actuator. The First analytic task merges the data streams from the light sensors and CCTV sensors to develop awareness on traffic congestion across the dependent road segments. The second analytic task combines output from first analytic tasks with appropriate traffic flow computational model to develop aggregated knowledge of traffic flow and congestion across the segments. The traffic

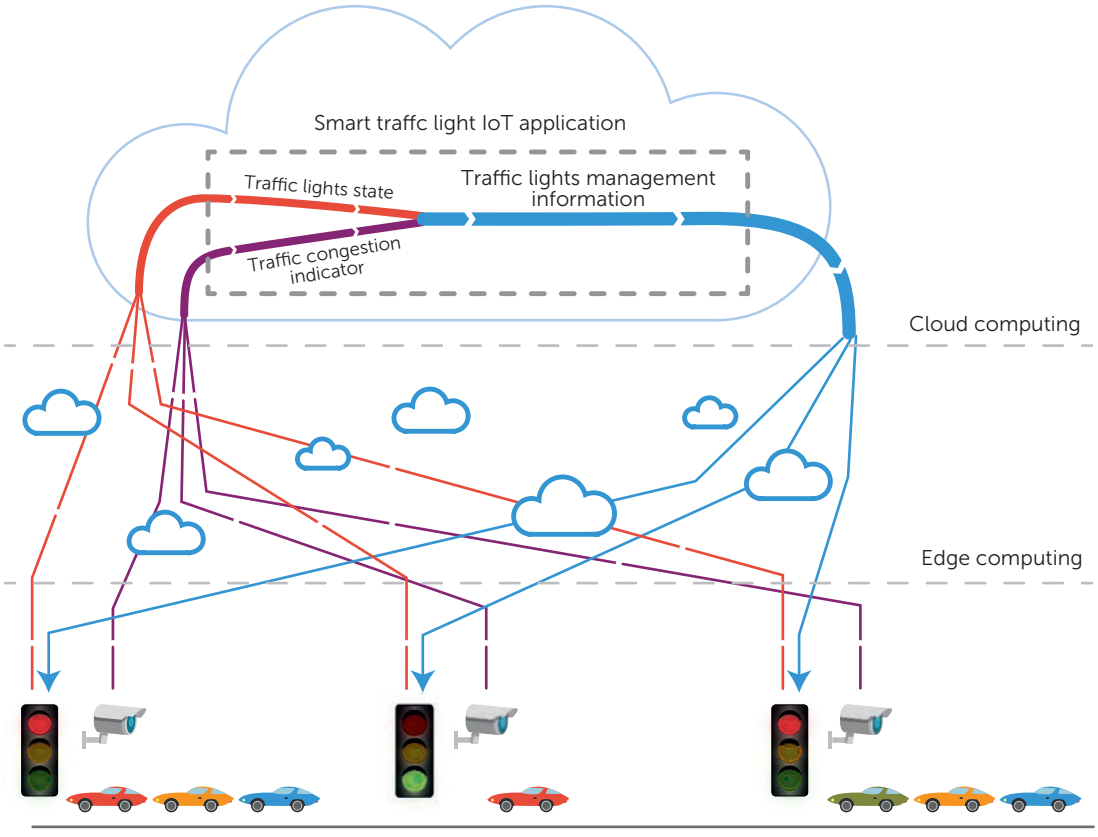


FIGURE 1. A high-level description of the smart traffic light IoT workflow application

flow computational model dynamically emits the traffic light control commands to the actuators about the switching off/on of the traffic lights across the segments such as that it leads to optimal traffic flow.

The classic approach for realising this kind of IoT workflow application (see Figure 1) relies exclusively on CDC resources, which could be distant from data sources, hence leading to excessive event detection (e.g., traffic congestion) delay. For example, the first analytic step of the traffic light sensor data aggregation should be mapped to nearby EDC resource, while resource-intensive second analytic task should be mapped to CDC resource, as it needs to execute complex traffic flow computational model.

## Design Goals of the Osmotic Flow model

We identify the key requirements that drive the design of Osmotic Flow programming model.

**Scalability and Elasticity.** Due to the huge amount of data that will be processed in a real-time fashion, scalability represent a key design requirement for IoT workflow applications. For example, a recent analysis of a (single) healthcare-related IoT workflow application (with 30 million of users) showed data flows up to 25,000 tuples per second.[10] The Osmotic Flow model should consider scalability and elasticity by design, so that applications can automatically grow and shrink based on data volume and velocity.

**Focus on data transformation.** The IoT application providers, who wants to realise a data transformation workflow, desires to focus only on data transformation, without spending too much time on management or configuration operations. Hence, the Osmotic Flow model should enable an easy deployment interface where data transformations should be easily

defined and, at the same time, each transformation should be seamlessly mapped to either EDC or CDC based on performance needs. As a consequence, the deployment process is transparently performed by the underlying run-time engine. However, the application providers can customize the framework behavior so to better address his/her specific performance needs.

**Efficient composition of data transformations.** The Osmotic Flow model should support the composition of cross-workflow data transformations and linking, so to easily realise complex workflows. To this end, Osmotic Flow should consider by design the possibility of composing data streams coming from multiple, public IoT devices and applications, thus promoting the principle of sharing and reusability. Our Osmotic Flow model should allow the application provider to easily define new streams, which extract high value information from raw data, without worrying about low level concerns related to their runtime execution, such as resource allocation, streams deployment, elasticity, and governance.

**Network Awareness.** The emerging IoT environment calls for strong network awareness. The Osmotic Flow model should not neglect the presence of communication delays while performing the deployment of data transformation tasks to CDC and/or EDC.

## Main Entities in the Osmotic Flow model

In Osmotic Flow, as depicted in Figure 2, includes one or more input endpoint that receives data from an external data sources and one or more output endpoints, which emit processed or transformed data towards sinks or other streams. A workflow is characterised by the following elements:

- one or more *data sources*: a data source is an entity, potentially external to the system, that continuously generates events or data. For example, a data source can be an IoT device emitting temperature measurements or traffic congestion conditions.
- one or more *sinks*: a sink is a final endpoint in the data transformation flow (or pipelne). Interested parties can subscribe to the sink for receiving notification information (e.g., traffic congestion information requested by drivers in Figure 1).

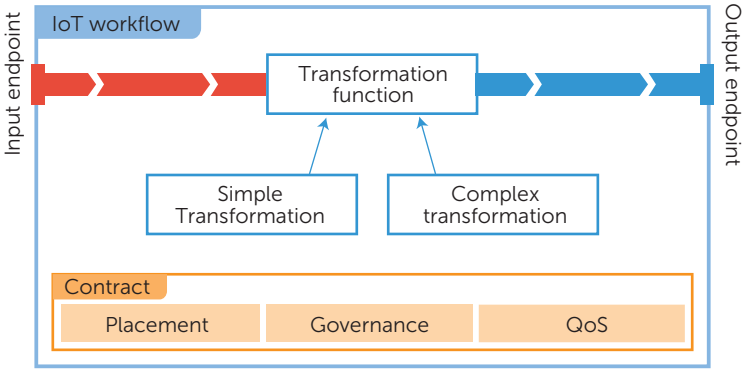


FIGURE 2. Depiction of IoT Transformation Functions in the Osmotic Flow model

- a *transformation function or task*: it encapsulates the user-defined analytics logic which transforms (e.g., combines, filters, splits) incoming data streams and passes the results to next transformation functions or final sinks of the workflow.
- a *contract*: it is a high-level configuration and performance requirement descriptor of the transformation functions or tasks.

## Types of IoT Data Streams

An IoT data stream can be ephemeral or public. An *ephemeral stream* is a special kind of stream that exists only if a sink is (directly or indirectly) interested to incoming flow. An indirect interest is manifest when one or more streams lie in between the stream and the final destination is the workflow endpoint (i.e., sink). Being ephemeral, the existence of the stream depends on the presence of (direct or indirect) interested sinks and its scope is restricted within the same application, i.e., it can be used only by user-defined transformations running within the same application that contains the stream. A public stream is a globally available stream and, as such, can be part of more than one IoT workflow application (for example traffic pattern stream data can be used for smart traffic management applications as well as air pollution monitoring applications).

## Types of Transformation Tasks

Transformation functions are the only piece of code that has to be defined by the application providers. A

transformation function encapsulates the data analytics logic. For an efficient execution, the stream model requires transformation functions to be as scalable as possible, therefore the latter are either stateless or provide an explicit definition of their state. We distinguish between two kind of transformations: simple and batch.

A *simple transformation* is applied to every incoming data in parallel and produces zero, one, or more outgoing data. For example, a simple transformation can be realised using only one type of data analysis programming model. An example of simple transformation (in context of Figure 1) could include tracking vehicles that exceed speed limit.

A *complex transformation* fuses one or more data streams before applying a transformation using one or more data analysis programming models (e.g., stream processing, batch processing, NoSQL). Complex transformation can produce zero, one, or more outgoing data streams. The group of incoming data streams fully determine the function state, which can then be manipulated by the transformation. A flow of data can be determined according to two modes: window and window-and-key. A window-based transformation creates a time-based or count-based window of events that have to be combined before running the transformation. For example, a window-based transformation can computes statistics on traffic patterns (see Figure 1) on a road segment in the last 30 seconds. On the other hand, a window-and-key transformation is a special case of windowed transformation that has a finer granularity in selecting the data streams for fusion. A classic example for a window-and-key transformation is the implementation of the vehicle counter task, which computes statistics on how many times a particular vehicle has travelled across a road-segment in last hour/week/month. Hence, the transformation across both historical and real-time data is dependent on the same key (i.e., vehicle registration number).

Simple transformations are stateless function, whereas complex transformations provide an explicit definition of their current state (real-time data) as well as it depends on the past state (historical data).

### Contract

The contract provides a high-level description of the IoT workflow application's configuration including:

- Placement constraints: these constraints guide the placement of transformation tasks over the distributed CDC+EDC infrastructure. If no restrictions are provided, the run-time execution engine can deploy a transformation task either on CDCs or on EDCs. Placement constraints be included to, e.g., maximize the utilization of nearby EDC resources or exploit centralized Cloud resources.
- Governance constraints: together with the previous one, the governance rules enable to specify further restrictions regarding the transformation task deployment and adaptation. These restrictions are often related to security, privacy, or law concerns. For example, a governance rule can exclude every edge resource belonging to a specific geographical region or can require to encrypt the exchanged data, so to meet stringent law restrictions.
- QoS or performance constraints: these ones express non-functional properties that should be met during the stream execution, so to obtain a desired quality level. For example, constrains can bound the maximum stream latency or minimum stream throughput or the event detection accuracy.

### Osmotic Flow Scheduling Architecture and Research Issues

Figure 3 provides a system-level description of the Osmotic Flow model. Whenever an IoT application provider wants to execute a workflow of data transformations, he/she submits the application code to the nearest Osmotic Resource Manager using a submission client. Then, the Osmotic Resource Manager allocates a new Node Manager, which, in turn, first determines the application placement, governance, and QoS constraints, and then distributes the data transformation tasks to appropriate EDC and/or CDC resources. The main software components include

**Osmotic Resource Manager (ORM).** An Osmotic Resource Manager coordinates Edge and Cloud resources and supports the Node Manager in determining the placement of Osmotic Flow transformation functions. In the proposed model, multiple

ORM can coexist and cooperate, where each one coordinates a pool of nearby CDC/EDC resources. The federation of multiple ORM enables the deployment of applications on the combined (EDC+CDC) infrastructure.

To ensure scalable communication and coordination between ORMs, future research should focus on developing self-healing load coordination protocols that can cope with changes in the infrastructures and IoT device state, and that can dynamically adapt to failures, connections, and detachments of ORMs and EDC/CDC resources. Another research thread could be to develop cooperative and opportunistic workload coordination protocol such that ORMs are able to balance their workload with each other in order to make sure that no CDC/EDC resource are wasted due to redundant data streams. Several workload coordination solutions already exist for CDC environments (e.g., Quincy, Omega, Sparrow, Mercury),[11] nevertheless the features of this new environment (CDC+EDC), as well as the characteristics of the Osmotic Flow model, foster the development of new load coordination policies and protocols, tailored for the specific setting where a significant heterogeneity of resources as well as multiple of data transformation tasks has to be management.

**Universal Stream Repository (USR).** To enable the sharing and reuse of high-value streams, Osmotic Flow includes a Universal Stream Repository (USR), which collects and provides the descriptor of every public stream available in the Osmotic Flow ecosystem. An IoT application provider relies on the descriptor to discovery existing streams and reuse them within his/her application.

Therefore the future research should focus on developing holistic data model for expressing EDC/CDC resources and data stream characteristics using an ontology-based representation, which enables encoding both dynamic (e.g., performance, status of stream) and static (e.g., functionality provided, types of events) QoS parameters. The ontology will thus guide decisions made on the types of data transformation tasks that are deemed most suitable for the deployment in the CDC or at the network Edge. An open access USR should be built. Existing ontologies such as Semanrtic Sensor Network (SSN) can describe concepts as IoT sensor characteristics, or data formats; nevertheless, they are not suitable to capture characteristics relevant to CDC/EDC resources.

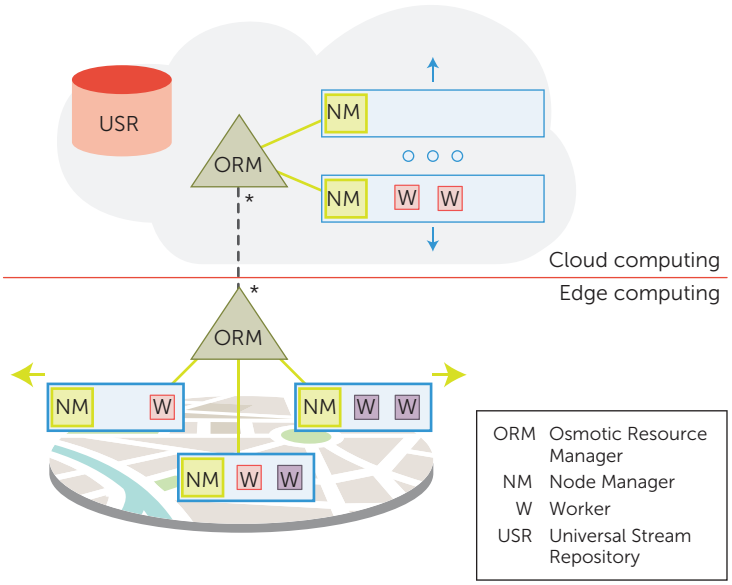


FIGURE 3. Scheduling Architecture of Osmotic Flow Model

**Node Manager.** The Node Manager is a per-machine agent that supports the ORM in controlling the available resources of the EDC/CDC infrastructure. Besides launching and terminating the execution of workers, the Node Manager monitors and reports statistical information about resource utilization (i.e., CPU, memory, network) to the ORM. Moreover, the Node Manager provides information to the ORM for determining the communication delays of the node with the other components of the infrastructure. Observe that communication delays can be obtained by means of either active/passive measurements (e.g., with a network coordinate system[12]), or with some network support (e.g., SDN).

Though Simple Network Management Protocol (SNMP), using the Management Information Base (MIB), has been highly adopted for monitoring resources in CDC, it lacks the ability to monitor EDC resources (as identified by the Internet Engineering Task Force[13]) due to huge computational overhead and the constrained nature of Edge resources. Hence,

future research will need to investigate modeling of a novel Edge resource monitoring agent that harnesses lightweight IoT protocols such as Constrained Application Protocol and a new Edge resource-specific MIB interface provided by the Internet Engineering Task Force.[13]

**Worker.** The Worker is in charge of executing one or more transformation functions. To this end, a worker collects data from the stream data source (i.e., another worker or an external data source), runs the user-defined transformation function, and emits outgoing streams. In other words, the worker takes care of the distributed execution of the user code, that defines only how to manipulate input data to obtain output data. Since streams are executed by workers, they directly communicate to transfer data up to the final consumers. Being stateless or with a window-based state, multiple transformation functions can run concurrently in a worker, and multiple workers can run concurrently on the infrastructure. Moreover, since a transformation function is defined with a fine granularity (i.e., per event or per window), it can be transparently scaled as the number of incoming events increases or decreases, up to—theoretically—creating an instance per each event.

As Osmotic Flow model thrives to support multiple type and mix of data transformation tasks on shared EDC+CDC infrastructures, the Worker need to be equipped with scheduling intelligence to automatically discover and resolve contention between co-deployed data transformation tasks. During deployment of data transformation tasks, the Worker must consider which data transformation tasks should be combined on an EDC and/or CDC resource, to minimize resource contention due to workload interference. Workload resource consumption and QoS are not additive, so understanding the nature of their composition is critical to deciding which transformation tasks can be deployed together. Existing content detection approaches such as Paragon[14] that applies collaborative filtering techniques for resolving contention between co-deployed, hypervisor-based application workloads on CDC are agnostic to the new hardware (e.g. Raspberry, Pi 3, UDOO board, Cisco IOx) and virtualisation features (e.g., Containers, Unikernels) of EDC resources.

Till data, several data analytic programming models and frameworks have been proposed. Nevertheless, most of them are designed to run in CDC, thus neglecting the presence of EDC resources. Osmotic Flow builds on the strengths of existing solutions and creates a novel approach for executing data analytics in a Cloud-supported Edge environment. Similar to Google Cloud Dataflow[8] and Apache Spark, Osmotic Flow defines a very simple and scalable programming model that enables to automatically deploy transformations with a high degree of parallelism. However, differently from existing approaches, Osmotic Flow focuses mainly on processing continuous and unbounded streams of data on decentralized CDC+EDC resources. Moreover, since the Osmotic Flow includes resource management capabilities, it can optimize how Edge and Cloud nodes are allocated among multiple and concurrent data transformation functions.

## References


1. Amazon Web Services, https://aws.amazon.com
2. F. Bonomi, R. Milito, J. Zhu, and S. Addepalli. Fog computing and its role in the Internet of Things. In *Proc. of MCC '12*, pages 13–16. ACM, 2012.
3. Google Edge Network, https://peering.google.com.
4. OpenFog Consortium, https://www.openfogconsortium.org.
5. M. Villari, M. Fazio, S. Dustdar, O. Rana and R. Ranjan, "Osmotic Computing: A New Paradigm for Edge/Cloud Integration," in *IEEE Cloud Computing*, vol. 3, no. 6, pp. 76-83, Nov.-Dec. 2016. doi: 10.1109/MCC.2016.124.
6. M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica. "Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing". In *Proc. of USENIX NSDI '12*, 2012.
7. S. Kulkarni, N. Bhagat, M. Fu, V. Kedigehalli, C. Kellogg, S. Mittal, J. M. Patel, K. Ramasamy, and S. Taneja. "Twitter Heron: Stream Processing at Scale". In *Proc. of SIGMOD '15*, 2015.
8. Google Cloud Dataflow, https://cloud.google.com/dataflow/
9. V. Gulisano, R. Jiménez-Peris, M. Patiño-Martínez, C. Soriente and P. Valduriez, "StreamCloud: An Elastic and Scalable Data Streaming

System," in *IEEE Transactions on Parallel and Distributed Systems*, vol. 23, no. 12, pp. 2351-2365, Dec. 2012.
10. R. Cortés, X. Bonnaire, O. Marin, and P. Sens, "Stream processing of healthcare sensor data: studying user traces to identify challenges from a big data perspective", *Procedia Computer Science*, vol. 52, 2015, pp. 1004-1009, 2015.
11. P. Pietzuch, J. Ledlie, J. Shneidman, M. Roussopoulos, M. Welsh, and M. Seltzer "Network-aware operator placement for stream-processing systems". In *Proc. of IEEE ICDE '06*, 2006.
12. F. Dabek, R. Cox, F. Kaashoek, and R. Morris. "Vivaldi: A decentralized network coordinate system". *SIGCOMM Comput. Commun.* Rev., 34(4), 2004.
13. P. V. D. Stok et al.,"CoAP Management Interfaces," October 2016, IETF Internet-Draft Work-in-Progress, https://datatracker.ietf.org/doc/draft-vanderstok-core-comi/.
14. C. Delimitrou and C. Kozyrakis. 2013. Paragon: QoS-aware scheduling for heterogeneous datacenters. *SIGPLAN Not.* 48, 4 (March 2013), 77-88. DOI=http://dx.doi.org/10.1145/2499368.2451125.
15. D. Puthal, S. Nepal, R. Ranjan and J. Chen, "Threats to Networking Cloud and Edge Datacenters in the Internet of Things", *IEEE Cloud Computing*, vol. 3, no. 3, pp. 64–71, 2016.

**MATTEO NARDELLI** *is a PhD student of computer science at the University of Rome Tor Vergata, Italy. His research interests are in the field of distributed computer systems, with a focus on the execution of data stream applications in geographically distributed environments. Contact him at nardelli@ing.uniroma2.it.*

**STEFAN NASTIC** *is a Postdoctoral Research Assistant at the Distributed Systems Group at TU Wien, Austria. His research interests include: Internet of Things and Edge Computing; Cloud Computing; Big Data Analytics; and Smart Cities. He received a PhD in software engineering and Internet technologies from TU Wien. Nastic has been involved in several EU-funded research projects such as SMART-FI, U-Test and SM4ALL, as well as, large industrial projects such as Pacific Controls Cloud Computing Lab (PC3L). Contact him at snastic@infosys.tuwien.ac.at.*

**SCHAHRAM DUSTDAR** *is a full professor of computer science heading the Distributed Systems Group at TU Wien, Austria. His work focuses on Internet technologies. He is an IEEE Fellow, a member of the Academy Europeana, and an ACM Distinguished Scientist. Contact him at dustdar@dsg.tuwien.ac.at or dsg.tuwien.ac.at.*

**MASSIMO VILLARI** *is an associate professor of computer science at the University of Messina. His research interests include cloud computing, Internet of Things, big data analytics, and security systems. Villari has a PhD in computer engineering from the University of Messina. He's a member of IEEE and IARIA boards. Contact him at mvillari@unime.it*

**RAJIV RANJAN** *is a reader in the School of Computing Science at Newcastle University, UK; chair professor in the School of Camputer, Chinese University of Geoscience, Wuhan, China; and a visiting scientist at Data61, CSIRO, Australia. His research interests include grid computing, peer-to-peer networks, cloud computing, Internet of Things, and big data analytics. Ranjan has a PhD in computer science and software engineering from the University of Melbourne (2009). Contact him at raj.ranjan@ncl.ac.uk or http://rajivranjan.net.*

*This article originally appeared in*
IEEE Cloud Computing, *vol. 4, no. 2, 2017.*

Editor: Yong Cui • cuiyong@tsinghua.edu.cn

# Internet of Things for Smart Cities: Interoperability and Open Data

**Bengt Ahlgren** • *SICS Swedish ICT*

**Markus Hidell** • *KTH Royal Institute of Technology, Sweden*

**Edith C.-H. Ngai** • *Uppsala University, Sweden*

The Internet of Things (IoT) for smart cities needs accessible open data and open systems, so that industries and citizens can develop new services and applications. As an example, the authors provide a case study of the GreenIoT platform in Uppsala, Sweden.

Today's cities face a variety of challenges, including job creation, economic growth, environmental sustainability, and social resilience. Emissions from motor vehicles have become a major source of air pollution in the world's large and medium-sized cities. Many large cities experience serious air pollution and greenhouse gas emission (GHG), which is made worse by increasing traffic congestion. With these challenges in mind, the European Union and many other countries are investing in information and communication technology (ICT) research and innovation, and developing policies to improve the quality of life of citizens and sustainability of cities. Given the trend of ICT for smart sustainable cities, understanding where we are in the evolution of the Internet is critical to future city-planning processes.

The Internet of Things (IoT) has been viewed as a promising technology with great potential for addressing many societal challenges. Cisco believes that many organizations are currently experiencing the IoT, the networked connection of physical objects and the cyberspace.[1] According to the International Data Corporation (IDC)'s *Worldwide Internet of Things Forecast, 2015–2020*, 30 billion connected (autonomous) things are predicted to be part of the IoT by 2020 (see www.idc.com/infographics/IoT). The IoT market size is forecast to grow from US$157 billion in 2016 to $661 billion by 2021.[2] The adoption of cloud platforms, development of cheaper and smarter sensors, and evolution of high-speed networks are expected to drive the growth of the IoT market.

Many cities, such as London and New York, see the increasing need and interest of the public sectors to explore IoT technologies to improve traffic flow, reduce pollution and energy consumption, and collect data for policing. Smart cities are an urban development vision to integrate multiple ICT solutions to manage a city's assets to create a sustainable environment, improve the quality of life, and enhance efficiency and economical value. The number of new IoT products and applications has grown exponentially in recent years. Various communication standards and protocols have been suggested in the community, and some have been adopted in different IoT devices. However, there are also quite a few proprietary protocols and cloud services in the IoT, which make the interoperability and sharing of data across different devices and platforms quite challenging. Open data in smart cities means not only global data collected and opened by the government, but also includes the sharing

 Published by the IEEE Computer Society 

Table 1. Standardized IP-based communication protocols for Internet of Things (IoT) devices.

| Layer | Protocol | |
|---|---|---|
| Application | IETF Constrained Application Protocol (CoAP)/REST engine | Message Queuing Telemetry Transport (MQTT) |
| Transport | UDP | TCP |
| Network | IPv6, RPL | |
| Adaptation | IPv6 over low-power wireless personal area networks (6LoWPAN) | |
| Media access control (MAC) | Carrier sense multiple access (CSMA) | |
| Physical | IEEE 802.15.4 | |

of data among individual citizens and industries with the government and general public. In this article, we'll discuss the advantages of open data and standards within the IoT, current limitations, and future trends.

## IoT for Smart Cities

The IoT provides individuals, society, and the business world new opportunities to access volumes of data and to develop new applications and services for creating a cleaner environment and more intelligent society.[3] The information society is rapidly becoming a central pillar for urban planners, architects, developers, and transportation providers, as well as in public service provision. One good example is using smartphones and smart meters to regulate energy consumption in the Hyllie smart networks of Malmö, Sweden.[4] The system enables people to measure, monitor, control, and influence their own energy consumption, and be able to independently produce renewable energy (for example, by using solar panels). One way to optimize the use of renewable energy and reduce costs is to decide how and when you want to charge your electric car. Consumers are informed of the supply of renewable energy in the system and how much electricity costs via smartphones or tablets.

From a public sector leadership perspective, cities can be viewed as microcosms of the interconnected networks for building a clean, energy-efficient, and sustainable society. In Amsterdam, a network-enabled LED streetlighting system has been developed to reduce the city's energy consumption and costs.[5] Similarly, in the US, Cisco and a wide range of public and private stakeholders in Chicago have been driving smart community initiatives to improve neighboring services and the quality of life.[6] IoT solutions are more effective when they facilitate open data and encourage public engagement, to achieve the goals of increasing productivity, decreasing costs, and improving citizens' quality of life.

## Interoperability and Open Standard Development

With the popularity of IoT devices, many IoT protocols and standards have been developed. In contrast to ordinary computers, IoT devices are normally constrained when it comes to memory space and processing capacity. In addition, IoT devices might be deployed where there's limited or no access to continuous power supply, which means that they need to operate under power supplied from batteries or small solar panels. As a consequence, power-efficient communication protocols with small memory footprints and limited demands on processing have been developed to support IoT devices. Traditional TCP/IP protocols haven't been designed with these requirements in mind. Over the past years, however, IoT protocols have been standardized on virtually all layers of the protocol stack. These protocols typically have low complexity as an important design goal and are optimized for constrained environments.

Table 1 shows a few examples of IP-based open protocol standards commonly used for IoT communication. For instance, IEEE 802.15.4 has been widely adopted in many smart devices as the MAC and Physical layer protocol. Several network layer and application layer protocols have also been proposed for constrained devices. Standard protocols are important to guarantee interoperability of different IoT devices.

However, using open standards doesn't automatically result in open systems. In our context, an open system means an integrated open IoT infrastructure solution for smart cities, providing access to open data and APIs for cloud services. In many cities, that infrastructure will be paid for, at least in part, by the city authorities using public funding. To motivate this investment, and get the most benefit for society, we argue that any smart city IoT infrastructure needs to be a truly open system, where equipment from many vendors can be used, and where the generated data can be more or less freely used by anyone to develop new services, based on low-level as well as processed sensor and IoT data. This kind of system will maximize innovation in the IoT domain, much as the Internet has done for information and communication services.

Many current IoT systems – for example, for air quality monitoring or the smart home – are either incomplete systems with limited functionalities (that is, in terms of sensing, storage, and analytics), or are closed, proprietary systems dedicated for a particular task. The latter are vertically integrated systems, sometimes called *stove pipes* or *vertical silos*, which can't be combined or extended

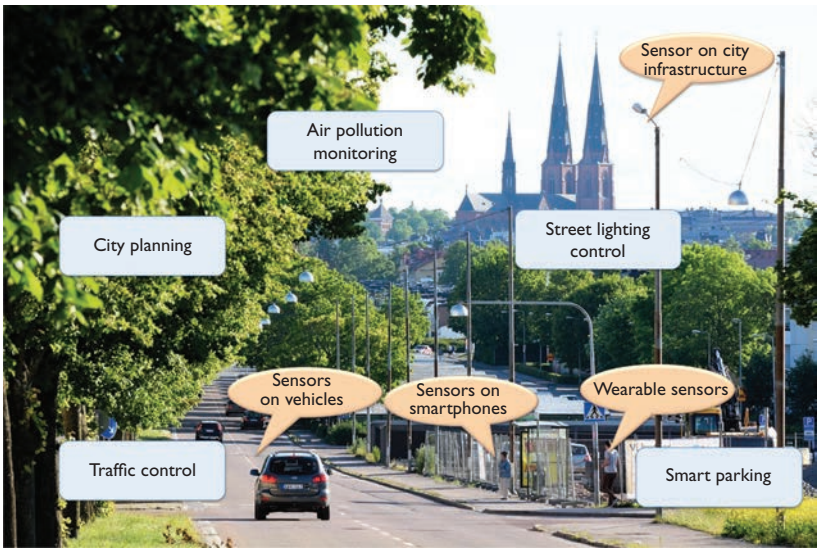


*Figure 1. An IoT system that includes heterogeneous sensors to collect data for smart city development. (Photo provided courtesy of Bengt Ahlgren.)*

easily with third-party components or services. The result is that once invested in a particular system, you're locked into that vendor's system. Vertically integrated systems are particularly problematic for the public sector, because this prevents fair competition in public procurement and is less suitable for large-scale data sharing.

Patrik Fältström[7] argues similarly that market forces work against open interoperability, especially in the IoT domain where, for example, a smart lighting system from one vendor only works with light bulbs from the same vendor. Systems are designed as end-to-cloud-to-end, where the cloud part is vendor-controlled with limited possibilities for third parties, and where the IoT devices often speak proprietary protocols to the cloud. Fältström argues that this lack of interoperability severely limits the market growth (for example, with smart light bulbs). Also, the dependence on a cloud service might render the device nonfunctional, should that cloud service for any reason, temporarily or permanently, disappear.

Instead of these stove pipes, we need horizontally designed systems with well-defined interfaces and data formats that can unleash the potential of open data, and that enable third parties to independently develop new applications and services, possibly combining several data sources. Providing open data has huge potential for innovation in digital applications and services, resulting in very large economic values. These interfaces (APIs) through which the IoT data can be accessed at multiple levels of refinement — from raw data directly from sensors, to highly processed data — also need standardization. The challenge is to provide an open system that lets users access the open data and cloud services without being locked by a particular platform. The open system should also allow third-parties to innovate based on the open data and open APIs.

## Case Study: GreenIoT Project in Sweden

We developed a GreenIoT solution that incorporates smart sensing and cloud computing technologies to encompass a more interactive and responsive city administration with private and public parties. The proposed open GreenIoT platform supports a wide range of applications, such as environmental monitoring, transportation, factory process optimization, and home security, and enables third-party innovation in new IoT-based services. Driven by Uppsala Municipality, we implement and demonstrate GreenIoT as a testbed in the city of Uppsala (the fourth largest city in Sweden) to support air pollution monitoring and traffic planning. Because the particulate level of Uppsala occasionally exceeds the EU standard, in particular during the winter and early spring, one objective is to reduce air pollution through active monitoring, traffic management, and better city planning.

Existing IoT technologies have largely contributed to hardware, software and protocol design. However, a major challenge of the IoT lies in how to extract valuable information from vast volumes of data generated from the smart devices (also known as the "Big Data" problem). Our GreenIoT solution leverages cloud computing to support intelligent data management, and integrate with green networking and sensing techniques to support energy-efficient and sustainable operations. The GreenIoT platform in Uppsala will be based on open standards, open to the public and supporting industries to test their new sensing products. It provides open data and open APIs for third parties to access the sensor data and make use of the cloud services. The open data generated by the smart devices and platform will drive the development of innovative applications and services.

One major goal of the project is an integrated solution for an environmental sensing system, which enables experimentation with applications and services using open environmental data, particularly for sustainable urban and transportation planning (see Figure 1). The GreenIoT architecture is manifested in terms of a testbed in Uppsala. The sensing system and application platform

are built from unique technology that provides open interfaces at several levels, energy and resource efficiency, and application independence. We use a unique tool for visualization in four dimensions, which supports smart city simulations and is fully integrated with the sensor data for real-time feedback. The testbed, including the open data and open APIs, allow third parties to develop and experiment new sensing products and services that could be exported to international markets.

To fulfill user requirements – from advanced tools for city planning as well as from novel applications making sensor data useful to citizens – we devised the GreenIoT architecture (see Figure 2).

Data produced by sensor networks are delivered through sensor gateways for storage and processing managed by cloud services for sensor data. The sensors use a publish/subscribe protocol, Message Queuing Telemetry Transport (MQTT), to communicate data in an open format through a broker for further storage and processing in the cloud, or for direct use by applications and services. We're also experimenting with information-centric networking[8] for direct access to sensor data.

Sensor data can be retrieved by tools and applications through well-defined APIs. The sensor data cloud services support both requests for raw sensor data and for pre-processed sensor data. Pre-processed data can be described as a grid of estimated values for a geographical region, where the values are calculated from the actual data produced by sensors in that region. A set of pre-processing types has been defined, such as interpolated data, hourly average, daily average, and weekly average. These types should be seen as a starting point, and more types are likely to be defined in the future. In the long run, it even should be possible for tools and applications to define processing that can be executed by the sensor data cloud services and then retrieve refined data according to their demands. The open APIs and open data format will facilitate the sharing of open data and guarantee the accessibility of cloud services without relying on a single device manufacturer or service provider.

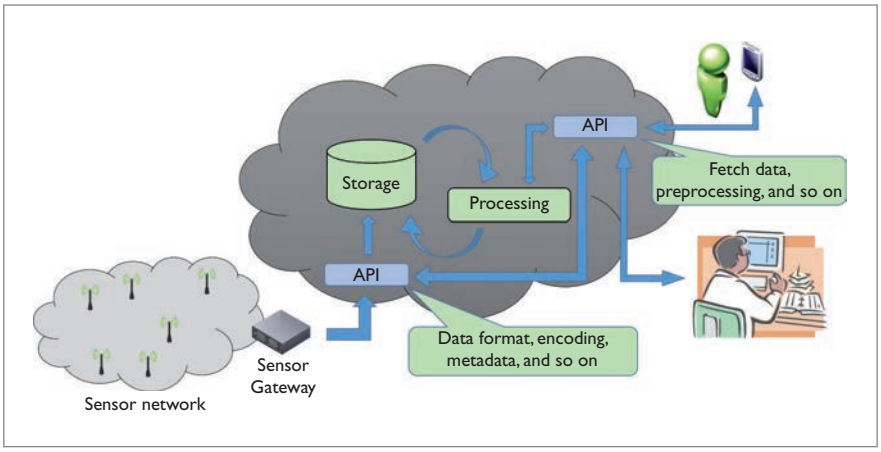


*Figure 2. The GreenIoT architecture. Our focus is on open access and interoperability, to fulfill user requirements – from advanced tools for city planning as well as from novel applications making sensor data useful to citizens.*

The vision of the "smart city," making use of the IoT to provide services for the good of citizens and public authorities, promises solutions to some of today's societal challenges such as air quality, transportation, and energy efficiency. These IoT systems must be based on open data and open standards, including protocols and interfaces, so that the systems enable third-party innovation in new services, and to avoid vendor lock-in. Standardized protocols might not be enough to achieve these goals – systems must be designed with openness in mind at all levels. Based on this concept, we designed and developed a GreenIoT platform in Sweden to demonstrate the benefits of open data and open platforms for smart city development. Over the next year, we will develop applications and carry out experiments using the Uppsala City IoT testbed, and formulate guidelines for public bodies for the procurement of open IoT infrastructure – including open APIs, common data formats, and how to avoid vendor lock-in. Open systems enabling innovation in new services are particularly important for publicly funded IoT infrastructures, to maximize the benefits for society.

**Acknowledgments**

This work is supported in part by the GreenIoT project grant (2015-00347) from VINNOVA, Sweden's innovation agency, and in part by EIT Digital in the ACTIVE project.

**References**

1. S. Mitchell et al., *The Internet of Everything for Cities*, Cisco, 2015; www.cisco.com/web/strategy/docs/gov/everything-for-cities.pdf.
2. Research and Markets, *Internet of Things (IoT) Market by Software Solution (Real-Time Streaming Analytics, Security, Data Management, Remote Monitoring, & Network Bandwidth Management), Platform, Service, Application Domain, and Region – Global Forecast to 2021*, tech. report, Apr. 2016; www.researchandmarkets.com/research/gsjxb5/internet_of.
3. C. Zhu et al., "Green Internet of Things for Smart World," *IEEE Access*, vol. 3, Nov. 2015, pp. 2151–2162.
4. Malmö Stad, *Climate-Smart Hyllie – Testing the Sustainable Solutions of the Future*, Swedish Energy Agency, 2013; http://malmo.se/download/18.760b3241144f4d60d3b69cd/1397120343885/Hyllie+klimatkontrakt_broschyr_EN_2013.pdf.

5. Philips, "Connected Lighting System," press release, 2014; www.newscenter.philips.com/main/standard/news/press/2014/20140327-philips-gives-workers-smartphone-control-of-office-lighting-with-groundbreaking-connected-lighting-system.wpd#.VL46kS5rNow.
6. City of Chicago, "Digital Roadmap to Improve Quality of Life," press release, Apr. 2015; www.cityofchicago.org/city/en/depts/mayor/press_room/press_releases/2013/september_2013/mayor_emanuel_releasescityofchicagosfirstevertechnologyplan.html.
7. P. Fältström, "Market-Driven Challenges to Open Internet Standards," *Global Commission on Internet Governance Paper Series*, Centre for International Governance Innovation (CIGI), paper series no. 33, May 2016; www.cigionline.org/publications/market-driven-challenges-open-internet-standards.
8. B. Ahlgren et al., "A Survey of Information-Centric Networking," *IEEE Comm.*, vol. 50, no. 7, 2012, pp. 1024–1049.

**Bengt Ahlgren** is a senior researcher in the Decisions, Networking, and Analytics (DNA) lab at SICS Swedish ICT. His current research focus is on designing networks based on an information-centric paradigm, where storage for caching is integrated in the network infrastructure. Ahlgren has a PhD in computer systems from Uppsala University, Sweden. Contact him at bengta@sics.se.

**Markus Hidell** is an associate professor in communication systems and at the Network Systems Laboratory (NSLab) at the KTH Royal Institute of Technology, Sweden. His current research interests are in the area of communication protocols and network architectures, including network virtualization, energy efficiency, and the Internet of Things (IoT). Hidell has a PhD in telecommunication from the KTH Royal Institute of Technology. Contact him at mahidell@kth.se.

**Edith C.-H. Ngai** is an associate professor in the Department of Information Technology, Uppsala University, Sweden, and a visiting researcher at Ericsson Research. Her research interests include the IoT, mobile cloud computing, information-centric networking, smart cities, and urban computing. Ngai has a PhD in computer science and engineering from the Chinese University of Hong Kong. She's a senior member of IEEE and a member of the ACM. Contact her at edith.ngai@it.uu.se.

*This article originally appeared in* IEEE Internet Computing, *vol. 20, no. 6, 2016.*

# Move Your Career Forward

## IEEE Computer Society Membership

## Explore These Internet of Things Resources


### *IEEE Pervasive Computing*

Designed for researchers, practitioners, and educators, *IEEE Pervasive Computing* explores the role of computing in the physical world as characterized by the Internet of Things, ubiquitous computing, and the like.

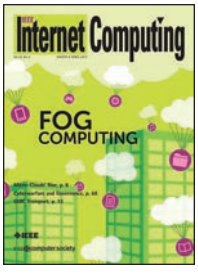


### *IEEE Internet Computing*

With insightful columns and peer-reviewed feature articles on all aspects of Internet computing, from programming and standards to security and networking, *IEEE Internet Computing* is the top choice for industry and academic readers.

### IEEE Computer Society Technical Committee on Very Large Scale Integration

The TCVLSI addresses the interactions among the various aspects of VLSI design including semiconductor processes and system-level, logic-level, and circuit-level design. It also covers computer-aided design techniques that facilitate the VLSI design process.

**FOR DIRECT LINKS TO THESE RESOURCES, VISIT**

www.computer.org/edge-resources

IEEE computer society

*The Community for Technology Leaders*

From the Editor in Chief

Editor in Chief: Maria R. Ebling ■ IBM T.J. Watson Research Center ■ ebling@us.ibm.com

# IoT: From Sports to Fashion and Everything In-Between

***Maria R. Ebling,*** *IBM T.J. Watson Research Center*

Pervasive computing and the Internet of Things have widespread impact across fields as seemingly diverse as sports and fashion.

## IOT IN SPORTS

It's easy to find IoT in sports, whether you're interested in wearables or objects. I'm sure the vast majority of *Pervasive Computing* readers own—or at least have seen—fitness bracelets such as the FitBit. Less common wearables include clothing such as the Hexoskin smart shirt, which measures your heart rate, heart-rate variability, breathing rate, and breathing volume in addition to your steps and pace (www.hexoskin.com). It can also track your sleep, measuring your heart rate, breathing, and sleep positions throughout the night.

This type of technology extends beyond sports to monitor people working in extreme and high-risk environments, such as firefighters and military personnel. Smart garments can send an alarm, including a location, if a firefighter or soldier is hurt in the line of duty.

Beyond wearables, Adidas now makes a smart soccer ball, called the miCoach smart ball, which has sensors embedded within the ball and retails for US $200. The sensors can detect the speed, spin, strike, and flight path and can send the data back to your smartphone via the miCoach app. The ball is a regulation size 5, and, amazingly, the battery life lasts for approximately 2,000 kicks over the course of a week (www.adidas.com/us/micoach-smart-ball/G83963.html).

IoT is also influencing golf. GolfTEC and K-VEST use sensors and display technology to provide feedback to golfers. The sensors in GolfTEC measure launch angles, spin rates, club speed, and the like (www.golftec.com/about-golftec/technology). Video allows the golfer to observe his or her swing from multiple angles. The K-VEST measures the golfer's hip, shoulder, and hands to provide feedback on body position (www.kandicomergolf.com/technology/k-vest).

We're seeing more and more IoT technology influencing sports. Right now, the focus is on data collection, but I anticipate such data will eventually be used to teach us how to play a new sport and optimize our performance and fitness levels.

## IOT IN FASHION

Perhaps more surprising are the inroads that IoT is making into the fashion industry. In the coming years, we can expect our clothing and accessories to come with RFID tags that retailers can use to manage their supply chain and reduce both theft and counterfeiting—examples include Avery Dennison (http://rfid.averydennison.com/en/home/solutions/apparel-and-retail.html) and Evrythng (https://evrythng.com/activate-digital-identities-for-products). In addition, consumers can also use these tags to find lost items and perform other functions, such as

- ensure we don't accidentally throw a hand-wash item into the washing machine or a line-dry item into the dryer,
- search for a duplicate or a new version of the same product, and
- learn fashion tips for how to wear the item or pair it with other clothes and accessories that we already own or could purchase.

I can even envision manufacturers and retailers someday using the tags to learn about the product's "end of life" in my closet, discovering why I'm donating the item or throwing it out. Is it because the item has simply been worn so many times that it has reached its end of life, or has my fashion sense matured? Or did the item shrink even though I followed the laundering instructions? Such data might help companies make manufacturing decisions to improve the quality and durability of their clothing, and retailers could note my modified fashion preferences.

Another application of pervasive technologies in the fashion industry is

**MISSION STATEMENT:** ***IEEE Pervasive Computing*** is a catalyst for advancing research and practice in mobile and ubiquitous computing. It is the premier publishing forum for peer-reviewed articles, industry news, surveys, and tutorials for a broad, multidisciplinary community.

 Published by the IEEE Computer Society 2469-7087/17/$33.00 © 2017 IEEE

Cisco's StyleMe Virtual Fashion Mirror.[1] This display technology lets consumers virtually try on clothing, create outfits, and view different colors, all without entering a dressing room. The display overlaps the customer's reflection with pictures of clothing so that consumers can see what the outfit might look like on their body. Similarly, Panasonic's makeup mirror lets consumers virtually try on different types of makeup, from false eyelashes to eyeshadow and blush (www.youtube.com/watch?v=JtwVVhvEwU8). Beyond the ability to see how you'd look with various beauty products, the mirror can also teach an inexperienced consumer how and where to apply the various products.

The use of IoT in the fashion industry is in the early days, but I foresee some fascinating applications in the years to come—much beyond the LED-enhanced outfits that we occasionally see stars wearing to gala events.[2]

## IN THIS ISSUE

One area that will be important for moving IoT and pervasive computing technology forward is the theme of this issue. Energy harvesting is critically important for certain types of IoT applications in both the sports and fashion domains. Without it, we risk forcing users to recharge the technology too often for usability, reducing long-term adoption.

In addition to our theme articles on energy harvesting, we also have several feature articles. As I mentioned earlier, IoT technology might eventually be used to optimize fitness levels, and along those lines, our first feature article is "Fitness Applications for Home-Based Training," by Iman Khaghani-Far, Svetlana Nikitina, Marcos Báez, Ekaterina A. Taran, and Fabio Casati. In this article, the authors survey fitness applications across a variety of platforms, including smartphone platforms (Android and iOS), desktop platforms (Windows and Mac), and console platforms (Nintendo and Xbox), looking explicitly at support for older adults who are frequently more isolated and less active and for whom exercising at home is an appealing option compared to traveling to a gym. The authors look at four design dimensions: interaction mechanisms, monitoring and sensing capabilities, coaching and tailoring features, and persuasion and motivation strategies.

In our second feature article, Suining He, Bo Ji, and S.-H. Gary Chan from The Hong Kong University of Science and Technology present "Chameleon: Survey-Free Updating of a Fingerprint Database for Indoor Localization." He, Ji, and Chan strive to make it easy to keep fingerprint databases up-to-date in spite of updated access points (APs). They make the observation that unaltered APs cluster a user's location whereas altered APs tend to disperse the user's location. With this observation, they can use the unaltered APs to localize the user and use the location to update the fingerprint database of the altered APs. They study this approach both on their university campus and in a major airport. If their results can be replicated, this approach could have major implications for the maintenance of indoor fingerprint databases.

Our third feature article, "Emerging Trust Implications of Data-Rich Systems," by Bran Knowles, identifies the challenges in building users' perception of trust in pervasive systems. In the article, Knowles uses a hypothetical fitness application, called the Fangle-Bangle, to illustrate the concepts. The six challenges identified include how such systems let users verify the data collected, how the data chains of these systems are mapped and presented to users, and how users can be given insight into the underlying algorithms without overwhelming them with technology. Understanding and addressing these components to trust is critical to the long-term acceptance and, consequently, the ultimate success of pervasive computing technologies.

On a related note, in our Pervasive Health department, Kelly Caine discusses the many ways we inadvertently give away private health data in

## EDITORIAL BOARD CHANGES

First, I'm pleased to announce that board member Robin Kravets will be taking on the role of Associate Editor in Chief. I look forward to her contributions, especially given her expertise in the areas of mobile computing and communications.

I'd also like to introduce two new board members, Florian Michahelles and Daqing Zhang.

Florian Michahelles heads the Web of Things research group at Siemens Corporate Technology. The Web of Things team borrows methods from the Semantic Web and investigates how machines, devices, and sensors can communicate and share data based on ontologies and semantics without requiring prior defined communication standards. The research group is being formed in Berkeley by international researchers who collaborate with leading US universities in industry-funded and publicly funded projects. Michahelles received his PhD from the Swiss Federal Institute of Technology (ETH) Zurich. Contact him at florian.michahelles@siemens.com.

Daqing Zhang is a Chair Professor at the School of EECS, Peking University, China and Vice Chair of the Pervasive Computing Federation of China. His research interests include context-aware computing, mobile computing, big data analytics, and pervasive elderly care. Zhang obtained his PhD from the University of Rome "La Sapienza." He is the associate editor for *ACM Transactions on Intelligent Systems and Technology* and *IEEE Transactions on Big Data*. Contact him at dqzhang@sei.pku.edu.cn.

her article "Privacy Is Healthy." She discusses the limitations of the Health Insurance Portability and Accountability Act, especially as it relates to non-clinical data collected by fitness-tracking apps and health-monitoring systems. She gives recommendations for actions that users and developers can take to improve their own and their users' privacy. This call to action is critically important for developers: if users misunderstand the privacy implications of the applications they use and discover that the health information they thought was private has been released, our industry will face a major crisis in terms of consumer trust, and adoption of our applications will greatly suffer.

In our Smartphones department, Yu-Chih Tung and Kang G. Shin present "ForcePhone: Software Lets Smartphones Sense Touch Force." This work describes a mechanism by which smartphones that don't have the force-sensing hardware (that is, all Android phones and lower-end Apple phones) can use software-based force-sensing. The idea behind this technology is to make it easier to use your phone one-handed. As someone who has been temporarily one-handed due to a broken shoulder, the idea is very appealing. The biggest issue I had, though, was in aiming and snapping photos—a use case that was, unfortunately, not included in their user study.

In this issue's Conferences department, Mateusz Mikusz, Sarah Clinch, and Sougata Sen provide an overview of MobiSys, which was held this past June in Singapore, and ASSET, the first Asian Students Symposium on Emerging Technologies. MobiSys included 31 papers and two impressive keynotes and covered topics ranging from smart environments, to sensing, to location awareness, to security and privacy. This year's best paper award went to Endri Bregu and his colleagues for their paper, entitled "Reactive Control of Autonomous Drones," about a new way of controlling drones, such that sensor readings only trigger controlling components if the sensor values have changed. If you were unable to attend the conference this year, I encourage you to read through the summary. It will make you aware of the many papers relevant to this community and will help you identify those of most interest to you.

Finally, our Notes from the Community department covers everything from interactive light to Pokémon Go. Of particular interest was the discussion of Project Jacquard, an effort between Google and Levi's to create a jacket for urban cyclists who want to interact with their phones while cycling. Multitouch sensors woven into the cuff of the jacket let the cyclists answer or place calls and use navigation apps.

The breadth of our field is amazing. The fact that pervasive and IoT technologies can be applied in such diverse areas as fashion and sports helps illustrate that fact. Someone once accused me of thinking that fashion and sports are the same field. In fact, I might be guilty of that—at least when you consider the use of IoT technology within these fields. One could argue that the cycling jacket proves my point.

### REFERENCES


1. L. Fretwell, "Cisco StyleMe Virtual Fashion Mirror," Cisco, Dec. 2011; www.cisco.com/c/dam/en_us/about/ac79/docs/retail/StyleMeEngagementOverview_120611FINAL.pdf.
2. P. Chi, "Met Gala 2016: Claire Danes's Glow-in-the-Dark Gown Upstaged a Red-Carpet Robot Army," *Vanity Fair*, 3 May 2016; www.vanityfair.com/style/2016/05/met-gala-2016-red-carpet.



**Maria R. Ebling** is a director at the IBM T.J. Watson Research Center. She manages a team building systems capable of supporting a Smarter Planet while not forgetting about the people who use such systems. Ebling received her PhD in computer science from Carnegie Mellon University. She's a member of the IBM Academy of Technology, a distinguished member of the ACM, and a senior member of IEEE. Contact her at ebling@us.ibm.com.



*This article originally appeared in* IEEE Pervasive Computing, *vol. 15, no. 4, 2016.*

# Cybersecurity and the Future

**Sven Dietrich,** City University of New York

*We need our computing systems to perform as intended, unaffected by adversaries big or small.*

*This article originally appeared in* Computer, *vol. 50, no. 4, 2017.*

As the world becomes increasingly interconnected, with more devices from our daily lives becoming part of the Internet of Things (IoT), one issue will certainly prevail: cybersecurity. We care about our security and privacy, and this undoubtedly won't change over the next 50 years. There will be shades of gray—based on cultural perceptions, needs, and trends—but privacy is part of human nature and affirmed as such in the Universal Declaration of Human Rights. We need our computing systems to perform as instructed or intended, unaffected by adversaries big or small.

Computing devices exist in our businesses, homes, pockets, bodies (in the form of medical devices), hospitals and medical offices, transportation systems, industrial production and energy generation systems, financial and payment systems, voting systems, and even in the dams that protect our land from rising sea levels. These devices all have one thing in common: they execute code on one or more silicon chips. Code can be flawed or vulnerable to exploits, which can compromise the aforementioned massively proliferated systems with various degrees of impact on our world and lives.

While we continue to promote secure coding practices for existing and upcoming programming languages and microchips and to educate stakeholders about better cybersecurity practices, we need to think about legacy problems with devices that can't easily be changed and form an integral part of a critical infrastructure. We're also creating secure ways of verifying, patching, and updating our devices to maintain their functionality, integrity, and health at both the software and hardware level, as well as ways to let devices communicate securely.

We'll face new challenges when quantum computers become commonly available. As current computational, algorithmic, and distributed techniques allow us to bypass, crack, and compromise security protocols and cryptographic primitives using conventional computing (for example, the recent Google-led effort to find collisions in the cryptographic hash function SHA-1), we'll face even tougher obstacles when quantum computing techniques threaten to undermine our security and privacy safeguards (for instance, allowing for the factoring of RSA moduli or breaking other public-key cryptography mechanisms). Quantum computing in the hands of a few powerful actors would create an even greater imbalance. The security challenges that could arise from a quantum computing environment itself are yet to be seen, compared to the conventional containment challenges that exist now and in some cases can be mitigated (by secure enclaves in recent Intel chips, for example).

Interconnections pose a continued risk. We're developing novel ways of connecting old and new devices, vulnerable or not. We're using new Internet architectures and next-generation networks and protocols, potentially using software-defined components, and looking to ameliorate underlying flaws by incorporating lessons learned from the past. We'll likely fix a few mistakes but also introduce new attack vectors that we can't imagine yet.

Harnessing vulnerable devices into maliciously behaving networks such as botnets or as a source of exfiltrating private or sensitive information (such as email or health and financial data) will continue as long as devices remain connected and always on, vulnerable to flaws, improperly configured due to human error, or intentionally sabotaged by insiders. Just as bots migrated from government and university lab systems to home computers and mobile devices, and distributed denial-of-service attacks turned IoT devices (such as cameras, DVRs, home automation, and broadband routers) into botnets, we can only imagine what the next target will be, from smartwatches to connected toothbrushes and pacemakers.

As we add more devices to our lives and connect them to the somewhat fragile Internet (by cable, Wi-Fi, or some method yet to be discovered), it's important to remember that information can flow both ways, maliciously or not.

**SVEN DIETRICH** is an associate professor in the Mathematics and Computer Science Department at the City University of New York's (CUNY's) John Jay College of Criminal Justice, is on the doctoral faculty of the Computer Science department at the CUNY Graduate Center, and is on the IEEE Cybersecurity Initiative Steering Committee. Contact him at spock@ieee.org.

PREFERRED PLUS

TRAINING & DEVELOPMENT

RESEARCH

BASIC

STUDENT

# New Membership Options for a Better Fit

**And a better match for your career goals.** Now IEEE Computer Society lets you choose your membership — and the benefits it provides — to fit your specific career needs. With four professional membership categories and one student package, you can select the precise industry resources, offered exclusively through the Computer Society, that will help you achieve your goals.

IEEE computer society

**Learn more at www.computer.org/membership.**

# Achieve your career goals with the fit that's right for you.

## Explore your options below.

| Select your membership | Preferred Plus | | Training & Development | | Research | | Basic | | Student |
|---|---|---|---|---|---|---|---|---|---|
| | **$60** *IEEE Member* | **$126** *Affiliate Member* | **$55** *IEEE Member* | **$115** *Affiliate Member* | **$55** *IEEE Member* | **$115** *Affiliate Member* | **$40** *IEEE Member* | **$99** *Affiliate Member* | **$8** *Does not include IEEE membership* |
| *Computer* magazine (12 digital issues)* | ✓ | | ✓ | | ✓ | | ✓ | | ✓ |
| *ComputingEdge* magazine (12 issues) | ✓ | | ✓ | | ✓ | | ✓ | | ✓ |
| Members-only discounts on conferences and events | ✓ | | ✓ | | ✓ | | ✓ | | ✓ |
| Members-only webinars | ✓ | | ✓ | | ✓ | | ✓ | | ✓ |
| Unlimited access to *Computing Now*, computer.org, and the new mobile-ready myCS | ✓ | | ✓ | | ✓ | | ✓ | | ✓ |
| Local chapter membership | ✓ | | ✓ | | ✓ | | ✓ | | ✓ |
| Skillsoft's Skillchoice™ Complete with 67,000+ books, videos, courses, practice exams and mentorship resources | ✓ | | ✓ | | | | | | ✓ |
| Books24x7 on-demand access to 15,000 technical and business resources | ✓ | | ✓ | | | | | | ✓ |
| Two complimentary Computer Society magazine subscriptions | ✓ | | | | ✓ | | | | |
| myComputer mobile app | *30 tokens* | | | | *30 tokens* | | | | *30 tokens* |
| Computer Society Digital Library | *12 FREE downloads* | | *Member pricing* | | *12 FREE downloads* | | *Member pricing* | | *Included* |
| Training webinars | *3 FREE webinars* | | *3 FREE webinars* | | *Member pricing* | | *Member pricing* | | *Member pricing* |
| Priority registration to Computer Society events | ✓ | | | | | | | | |
| Right to vote and hold office | ✓ | | ✓ | | ✓ | | ✓ | | |
| One-time 20% Computer Society online store discount | ✓ | | | | | | | | |

** Print publications are available for an additional fee. See catalog for details.*

**www.computer.org/membership**

IEEE computer society

Editor: André Stork

# The Need to Help Journalists with Data and Information Visualization

**Susan Reilly**
*Florida Atlantic University*

As of 2015, 89 percent of mobile users (144 million users) accessed news via their mobile devices, and it is estimated that by 2020 mobile devices will account for two-thirds of all online activity.[1] In turn, smaller screen sizes reduce reading time and ease of reading, and mobile access reduces news seeking as well as interaction with graphical elements.[2] At the same time, the quantity and availability of news seems to be resulting in people reading shorter stories more frequently throughout the day.

As news migrates to mobile phones, media companies are turning to data visualization to whet readers' appetites for stories they can read at length later on their home or work computers. However, journalists are trained to write stories, not in statistics or coding. The big news organizations have the funds to hire computer graphics experts, but local news organizations need help.

## Why Do Journalists Need Data Visualization?

Data visualization has been around since antiquity with star charts and navigational maps. During the 19th century, with the growth of industrialization, large bodies of data were collected by businesses and governments to aid in product development and social planning. With the growth of data collections, ever-more sophisticated mathematical techniques have been used to analyze them. Governments and corporations continue to collect and store huge amounts of data. The introduction of human-computer interaction in the 20th century aided in the statistical analyses of these vast databases and the visual presentation of patterns, relationships, and hierarchies.[3]

The latest field to utilize data visualization is journalism. News organizations' enthusiasm is being driven by the relentless 24/7 Internet news cycle, which requires the presentation of fresh stories around the clock in order to stay competitive. As people become increasingly computer literate, it is hoped that online news readers will be drawn in by interesting data visualizations on their mobile devices and then later read the stories that were the motivation for those data visualizations on their home computers.

In a news environment where high story turnover is necessary, the strength of data visualization lies in the viewer's ability to process visual information more rapidly than verbal information. What seems to work is the clear presentation of relationships, the accurate representation of quantities, the easy comparison of qualities, and the clear ranked order of values.[4]

Yet, journalists are trained to work with words, using them to construct stories that explain events. Now they are being asked to add a visual component to those stories as well. Quite frankly, many feel overwhelmed. A professional journalist described her beat this way:

> Now my job isn't just to report, but also to find an audience for the story by doing social media promotion—Twitter and Facebook. Then I do format and production work. We have a content management system for print. I add a cutline and attach a photo. We have a free site and a paid site, so I produce a story for the free site that has a taste and a summary and then I produce a story for print that is more substantial. Now I have to think of [mobile] because six out of 10 readers are reading on their cell phones. We learned how to do all these fancy interactive graphics for the computer, but we found out that they don't work on phones. You need really simple graphics on phones.

 Published by the IEEE Computer Society 

## What Are Journalists Trying to Do?

Alberto Cairo, the Knight Chair in Visual Journalism at the University of Miami, has suggested that there are two ways for data visualization to be effective: one is an infographic that provides "spontaneous insight" (the aha moment) and the other is an infographic that provides "knowledge building insight" gained through the exploration of a complex system. The former can be understood immediately, while the latter takes time to explore.[5] The aha moment might work best on mobile devices that readers look at for short periods, multiple times during the day; the "knowledge building insight" might work better on home or work computers where the increased screen size facilitates longer periods of interactivity.

An example of a data visualization that creates an aha moment is Bill Bunge's iconic map, "Where Commuters Run Over Black Children on the Pointes-Downtown Track" (see civic.mit.edu/blog/kanarinka/the-detroit-geographic-expedition-and-institute-a-case-study-in-civic-mapping). Bunge's map was originally created by the Detroit Geographic Expedition, which addressed racial inequality in Detroit's inner core in the 1970s. The data was collected from police reports. With a shockingly simple legend, the map shows how commuter traffic from affluent white suburbs speeding through a black neighborhood repeatedly resulted in injuries to young children crossing the streets.[6] The simple visual representation of black children hit by white commuters' cars became an allegory for racial conflict in 1970s Detroit. Bunge's map is simple enough to be viewed easily on a mobile device and it provides, in Alberto Cairo's terms, spontaneous insight. This map is one of many Bunge constructed to visualize the findings of the Detroit Geographic Expedition. Yet, to get the full story, you still have to read the report.

An example of a data visualization providing Cairo's knowledge building insight can be found in the 2016 Pulitzer Prize winning investigative report "Failure Factories" by the *Tampa Bay Times* (www.tampabay.com/projects/2015/investigations/pinellas-failure-factories/chart-failing-black-students/). A team of journalists combined data visualization and reporting into a powerful series on the abandonment of school integration in Pinellas County, Florida. The reporters retrieved data from test scores, poverty rates, racial demographics, and school board meeting transcripts that showed the gradual transition of average schools into the worst in the state.[7] The journalists wrote the story, and the data director and data reporter created a data visualization that is so simple it works on mobile devices. But to read the full story, you have to use a desktop computer.

The first screen in the data visualization asks, "Why is Pinellas County the worst place in Florida to be black and go to public school?" Subsequent screens illustrate how after a decision by the Pinellas County School Board in 2007 to stop integrating schools, the percentage of black students began to rise in five elementary schools in black communities. At the same time, funding for those schools fell and teacher turnover increased. Ultimately, 95 percent of students at those five elementary schools failed the state's reading and math tests, the greatest concentration of academic failure in Florida. The final screen in the data visualization asks, "Who is responsible?" The investigative report resulted in the US Department of Education opening a civil rights investigation into whether the school district systematically discriminates against black children.

***Many journalists are trying to create data visualizations with little training and with programs that were developed for other professions.***

Large national news organizations, like the *New York Times* (with a circulation of more than 4 million, print and digital), have full-time employees trained in statistical methods, data mining, coding, and computer graphic applications who are tremendously helpful to the reporters investigating complex problems. By cleverly manipulating large numbers of records, a good statistician can locate the place where the story lies, and a good computer graphics expert can use a program like the D3.js JavaScript library to make it visible.

Many small local newspapers and television stations can't afford statisticians or computer graphics professionals, however. These organizations are making do by redefining the jobs of the existing journalists who are the most computer savvy and hoping for the best. Often a data visualization that works on a home computer is hard to read on a small screen or doesn't work in newsprint where the legends are in shades of gray. Badly designed data visualizations can be baffling and, worse, a waste of both the journalist's and the reader's time.

## What Do Journalists Need?

There are roughly 83,000 professional journalists working in newspaper and television newsrooms across the United States.[8] Many are trying to create data visualizations with little training and with programs that were developed for other professions. They spend too much time trying to adapt these programs to their needs. For example, the designated journalists/data analysts at two local newspapers in South Florida (with circulations of roughly 100,000) use Excel, Access, MySQL, R, and QGIS. They have to look at the numbers several ways before deciding the best way to tell the story in graphics. Since coding is problematic, they often rely on D3.js alternatives, like DataWrapper, Google Chart Tools, and StoryMap JS.

Local journalists need tools that make it easier to take raw data and convert them into interesting graphics without requiring advanced coding skills. It would be helpful if computer graphics researchers and designers would create a tool that featured a menu of options with recommendations to both organize and present visual information in ways that news consumers can understand quickly. This might include different ways to analyze statistics as well as suggestions for data-driven documents. Perhaps a good existing tool could be adapted, but it should be menu-driven rather than code-driven, and the menu should be built by journalists. The more intuitive the tool, the more journalists will be able to take advantage of data visualization.

Another key requirement in this field is responsive design for different screen sizes. News organizations need a way to customize headlines as readers move between devices. Some graphics don't scale down well, so journalists need a way to easily adapt static graphics to different screen sizes. A tool that allowed readers to tag small screen stories to read at length on their PCs would be ideal.

Without a strong business model to replace the revenue lost to Internet advertising, many in the journalism profession fear the continued disappearance of local news organizations. Despite the difficulty of finding ways to make it cost effective, local news is still something that communities value. Local news contributes to social cohesion—something that many feel is at risk as people become increasingly involved in their respective Internet worlds.

Hopefully, a dialog can begin between journalists and computer scientists about how to work together to help local journalists provide the information that citizens need in order to make good decisions in a democratic society.

## References


1. Knight Foundation, "Mobile-First News: How People Use Smartphones to Access Information," 11 May 2016; kf-site-production.s3.amazonaws.com/publications/pdfs/000/000/187/original/Topos_KF_Mobile-Report_Final_052616.pdf.
2. J. Dunaway, "Mobile vs. Computer: Implications for News Audiences and Outlets," Shorenstein Center on Media, Politics, and Public Policy, 30 Aug. 2016; shorensteincenter.org/mobile-vs-computer-news-audiences-and-outlets.
3. M. Friendly, "A Brief History of Data Visualization," *Handbook of Data Visualization*, C. Chen, W. Hardle, and A. Unwin, eds., Springer, 2008, pp. 15–56.
4. S. Few, "Data Visualization for Human Perception," *The Encyclopedia of Human-Computer Interaction*, 2nd ed., M. Soegaard and D. Rikke Friis, eds., Interaction Design Foundation, 2014.
5. G. McGhee, "The 'Rules' of Data Visualization Get and Update," *National Geographic*, 16 Oct. 2015; news.nationalgeographic.com/2015/10/151016-data-points-alberto-cairo-interview.
6. C. D'Ignazio, "The Detroit Geographic Expedition and Institute: A Case Study in Civic Mapping," 2013; civic.mit.edu/blog/kanarinka/the-detroit-geographic-expedition-and-institute-a-case-study-in-civic-mapping.
7. S. Nesmith, "Tampa Bay Times' Investigation Is a Model for How to Report on School Resegregation," *Columbia Journalism Rev.*, 19 Aug. 2015; www.cjr.org/united_states_project/tampa_bay_times_school_resegregation.php.
8. L. Willnat and D. Weaver, "The American Journalist in the Digital Age: Key Findings," School of Journalism, Indiana Univ., 2014; news.indiana.edu/releases/iu/2014/05/2013-american-journalist-key-findings.pdf.



***Susan Reilly*** *is a full professor of multimedia studies at Florida Atlantic University. She is a writer and producer of public television documentaries and the author of several books on critical media theory. Contact her at sreilly@fau.edu.*

*Contact department editor André Stork at andre.stork@igd.fraunhofer.de.*



*This article originally appeared in*
IEEE Computer Graphics and Applications, *vol. 37, no. 2, 2017.*

# Mutual Dependence Demands Mutual Sharing

**Daniel E. Geer Jr.**
In-Q-Tel

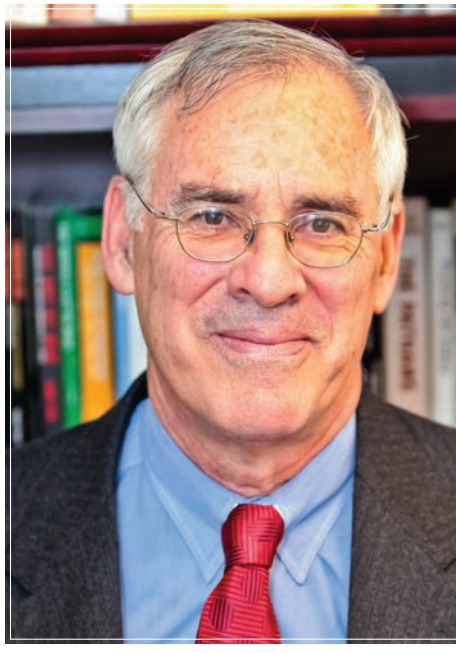

**Richard Danzig**
Johns Hopkins University

Unfortunately, interdependence, cooperation, and trust are poorly correlated. This is the fundamental axiom of cyberinsecurity risks for all of us.

Interdependence makes risk transitive; if A depends on B to function and B depends on C to function, then a failure of either B or C induces a failure for A. By contrast, trust is not transitive; that A trusts B and B trusts C does not impel A to trust C, nor would B's loss of faith in C impel A to lose faith in B. Worse, the fact that A vitally depends on B and C doesn't necessarily induce cooperation with either.

In the digital world, our security is as connected as our devices. Your firewall might be strong and your defenses might be active and robust, but if a sophisticated and determined opponent attacks your counterparties, spoofs your legitimate suppliers or customers, or infects your security providers, then that opponent will find a way to undermine your security. Most users don't even discover their penetration on their own—they learn it from others.[1]

In other words, not only is the problem technically and mathematically challenging, but its operational impact has a social component that compounds the complexity; you are probably insecure and you are dependent on others to tell you that.

According to a *Wall Street Journal* poll of its CEO Council: "9% of the respondents said they would never sufficiently trust the [US] government with information to work with it during a cyberattack ... . Another 34% said they would cooperate with the government only if their own company was being attacked."[2] We don't think an obligation to share information is simple. Companies have competitive concerns and face legal vulnerabilities when they acknowledge problems. Individuals and companies rightly value their privacy and fear overreaching by an overly intrusive government. But we don't think it's viable for nearly half of all CEOs, even in a small sample, to embrace non-cooperation with public authorities.

History shows that nothing unites like a common enemy, and that failure to unite guarantees defeat. Our shared security requires shared defenses.

We need a broad reporting requirement for cyberattacks and increased information sharing among government, private sector actors like ISPs, and private sector attack targets. This imperative can be seen in other contexts. Accident data must be reported when airplanes collide and workers are injured. Beginning in 1912, a diagnosis of plague, cholera, yellow fever, typhus, or smallpox obligated the doctor or clinic to share that information (by telegraph) with public health officials. That list of five communicable diseases has grown to 40, but the same rule applies: prompt, exceptionless disclosure for a limited set of priority conditions irrespective of privacy rules. The principle? Transitive risk above some threshold necessitates information sharing.

The US is stumbling toward coerced information sharing about digital attacks. We do this with different rationales in different contexts, such as when we require prompt and detailed attack information from defense contractors to Pentagon authorities, when state laws force disclosure if customers' credit card or other personal information is exposed, and when the SEC requires the announcement of security breaches that materially impair corporate operations. But this is piecemeal improvement, and we need to move beyond islands of insight. We should take advantage of opportunities presented to us by digital systems for immediate and comprehensive reporting. To benefit fully from the technical opportunity, we need to complement it with a mandatory reporting system that is comprehensive, inexpensive, adequately protective of confidentiality, and valuably informative about the volume, pattern, and character of digital attacks.

2469-7087/17/$33.00 © 2017 IEEE Published by the IEEE Computer Society

Voluntary systems fill some of the gaps. ISPs and equipment vendors capture attack data and alert their users with varying consistency, speed, and detail. Federal and private Information Sharing and Analysis Centers (ISACs) encourage exchanges between industrial sectors. Utilities representing perhaps two-thirds of America's customers participate in an automated Cybersecurity Risk Information Sharing Program (CRISP). Companies share information with chosen others, sometimes outside of their own industries to diminish competitive and antitrust concerns. Former colleagues commonly exchange information across industry boundaries.

However, these don't suffice. Large vulnerabilities remain when a third of the nation's utilities don't yet participate in CRISP. In a connected grid, vulnerabilities for some have consequences for all. Though industry-specific data compilation and analysis are commendable, software and hardware attacks cross industries and exploit common vulnerabilities in ways we are not well positioned to understand, never mind remediate.[3] Two of the most recent and insightful assessments have recommended improved voluntary reporting.[4,5] But why not mandatory?

For its part, the federal government needs to share more systematically. Often the FBI reveals to a surprised company that it has been penetrated, and then doesn't divulge any more about what it knows or how. Our intelligence establishment understandably guards its secrets, but its highest priority (to protect us) is undervalued in an effort to protect its sources and methods.

We don't think a regime of broader, compelled sharing is an easy one or without side effects. There is a great deal of thinking still to be done. Biological infections are treated by professionals, but malware infections are treated by amateurs. Diseases spread within jurisdictions before they become global, but malware is global from the get-go. Transmissible diseases can mutate, but they have studied, predictable behaviors, whereas malware comes from sentient opponents who can be intentionally devious.

Of course, our proposal is not the only way to proceed, but one can't ignore that its cousins have been proven successful in life-and-death matters. We believe the digital domain generates more widely shared vulnerabilities than any other. When dependence is mutual—that is to say when for A to function B must be operational and for B to function A must be operational—the brittleness of the combined system to cascade failure requires other coercive measures if either A or B is critically important to others beyond their mutual embrace.[6] ■

## References


1. Data Breach Investigations Report, Verizon, 2016; www.verizonenterprise.com/verizon-insights-lab/dbir/2016.
2. A. Cullison, "NSA Chief: 'Uneven Cooperation between Public, Private Sectors Impedes Cyber Defenses," *Wall Street J.*, 15 Nov. 2016; blogs.wsj.com/washwire/2016/11/15/nsa-chief-uneven-cooperation-between-companies-government-impedes-cyber-defenses.
3. P. Behr, "DOE Seeks to Offer Cyberthreat-Sharing Defenses to Small Utilities," E&E News, 6 Jul. 2016; www.eenews.net/stories/1060039828.
4. "From Awareness to Action: Cybersecurity Agenda for the 45th President," CSIS, 4 Jan. 2017; www.csis.org/analysis/awareness-action.
5. "Report on Securing and Growing the Digital Economy," Presidential Commission on Enhancing Nat'l Cybersecurity, 1 Dec. 2016; www.whitehouse.gov/sites/default/files/docs/cybersecurity_report.pdf.
6. D. Geer, "For Good Measure: Stress Analysis," ;login:, vol. 39, no. 6, 2014, pp. 56–57.



**Daniel E. Geer Jr.** is the Chief Information Security Officer of In-Q-Tel. Contact him at dan@geer.org.

**Richard Danzig** is a senior advisor to the Johns Hopkins University Applied Physics Laboratory and former Secretary of the Navy. Contact him at rjdanzig@gmail.com.



*This article originally appeared in*
IEEE Security & Privacy, *vol. 15, no. 1, 2017.*

# Careers Related to the Internet of Things

Florian Michahelles has run Siemens' Web of Things research group—which investigates the application of Semantic Web technologies to the Internet of Things (IoT)—since 2013. Having worked in the fields of ubiquitous and wearable computing for more than a decade, Michahelles' current focus at Siemens is leveraging Web and semantic technologies to enable new business opportunities, particularly in the fields of wearable sensing and human-robot interaction. He wrote "Internet of Things Reality Check" in *IEEE Pervasive Computing*'s April—June 2017 issue. We asked Michahelles about IoT-related careers.

***ComputingEdge:*** Which IoT-related careers will see the most growth in the next several years?

**Michahelles:** Any career bridging the disciplines of mechanical engineering, electrical engineering, design, computer science, interactive design, and communications will be in high demand because IoT reaches across these disciplines.

***ComputingEdge:*** What would you tell college students to give them an advantage over the competition?

**Michahelles:** Go beyond your major and think about also taking non-tech majors, such as by combining computer science and psychology, business and electrical engineering, or material science and sensors.

***ComputingEdge:*** What should applicants keep in mind when applying for IoT-related jobs?

**Michahelles:** Be an expert in one topic. While breadth is welcome, depth in one topic is key. Breadth then helps you effectively apply your expertise.

***ComputingEdge:*** How can new hires make the strongest impression in a new position from the beginning?

**Michahelles:** Listen and learn, get your hands dirty, be bold and courageous in proposing new

2469-7087/17/$33.00 © 2017 IEEE Published by the IEEE Computer Society July 2017 

ideas. Play with technologies you haven't used before, and quickly build demos and prototypes to convey your ideas to others.

***ComputingEdge:*** Name one critical mistake that young graduates should avoid when starting their careers.

**Michahelles:** Don't be afraid of failing. Instead, be brave enough to fail often, but avoid failing twice at the same thing. Keep improving.

***ComputingEdge:*** Do you have any learning experiences that could benefit those just starting out in their careers?

**Michahelles:** First, find your passion and develop it. Passion is the prerequisite to being successful at something. Second, learn how to deal with people. How do you present your ideas? How do you explain your ideas to others? These days, it's really hard to create something innovative all by yourself. Therefore, it's important to learn to work with others. And third, get a sense of what is required. Find out what's needed, where the opportunities are, and adjust your passion to this need.

*ComputingEdge*'s Lori Cameron interviewed Michahelles for this article. Contact her at l.cameron@computer.org if you would like to contribute to a future *ComputingEdge* article on computing careers. Contact Michahelles at florian.michahelles@siemens.com.

IEEE computer society

## SkillChoice™ Complete

Now with expanded libraries and an upgraded platform!

Valued at **$3,300!**

**3,000+** online courses

**MENTORSHIP**

**28,000+** BOOKS

**6,000+** VIDEOS

**Practice Exams**

**15,000+** *Books24x7 titles*

**OVER 20x** as many resources as before

# One membership. **Unlimited knowledge.**

Did you know IEEE Computer Society membership comes with access to a high-quality, interactive suite of professional development resources, available 24/7?

Powered by Skillsoft, the SkillChoice™ Complete library contains more than $3,000 worth of industry-leading online courses, books, videos, mentoring tools and exam prep. Best of all, you get it for the one low price of your Preferred Plus, Training & Development, or Student membership package. There's something for everyone, from beginners to advanced IT professionals to business leaders and managers.

The IT industry is constantly evolving. Don't be left behind. Join the IEEE Computer Society today, and gain access to the tools you need to stay on top of the latest trends and standards.

Learn more at **www.computer.org/join**.

ACCESS TO SKILLSOFT IS AVAILABLE WITH