IQT
QUARTERLY
VOL. 8 NO. 1 SUMMER 2016

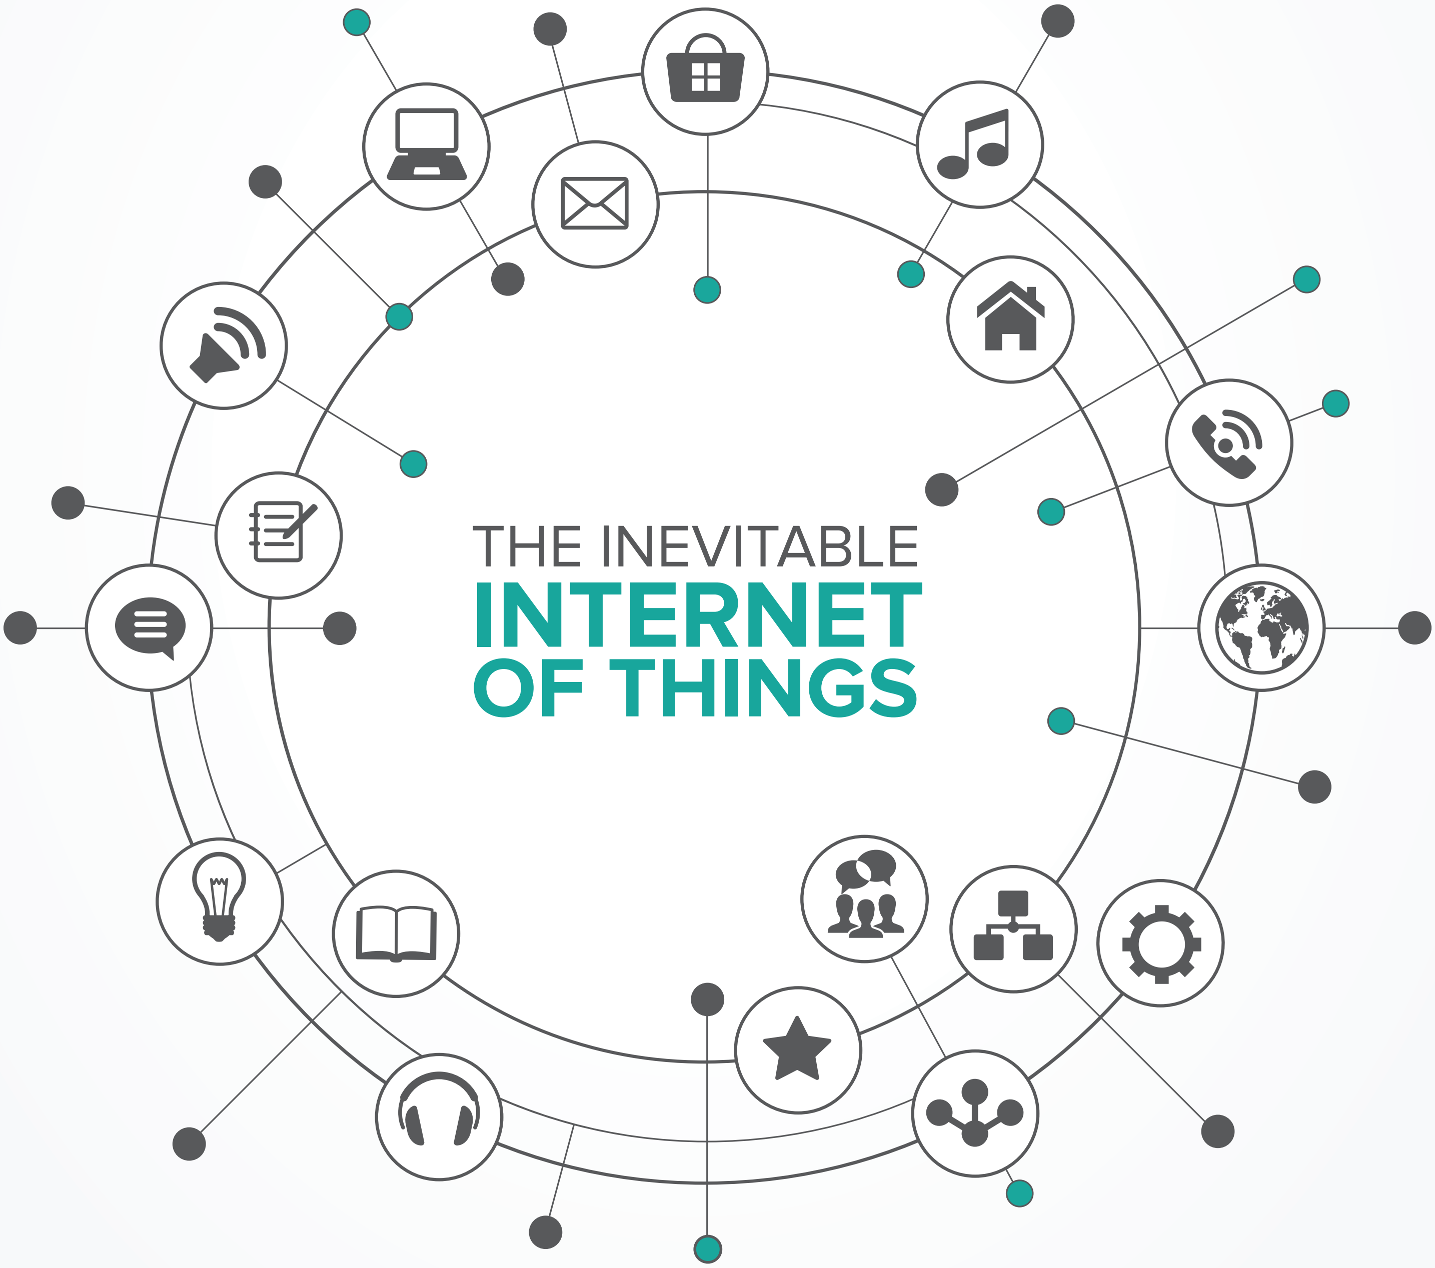

THE INEVITABLE INTERNET OF THINGS

IQT®
IN-Q-TEL

IQT Quarterly is a publication of In-Q-Tel, Inc., the strategic investment organization that serves as a bridge between the U.S. Intelligence Community and venture-backed startup firms on the leading edge of technological innovation. *IQT Quarterly* advances the situational awareness component of the IQT mission, serving as a platform to debut, discuss, and debate issues of innovation in the areas of overlap between commercial potential and U.S. Intelligence Community needs. For comments or questions regarding IQT or this document, please visit www.iqt.org, write to iqtquarterly@iqt.org, or call 703-248-3000. The views expressed are those of the authors in their personal capacities and do not necessarily reflect the opinion of IQT, their employers, or the Government.

EDITORIAL

IQT Quarterly, published by In-Q-Tel, Inc.
Editor-in-Chief: Adam Dove
Theme Editor: Ravi Pappu
Contributing Editors: Carrie Sessine, Brian Spatocco, Melissa Hayes. and Andrea Brennen
Printed in the United States of America

CONTENTS

VOL. 8 NO. 1

02
Making IRL Computable: The Inevitable Internet of Things
by Ravi Pappu

07
A Look Inside the Issue

08
The Genesis and Evolution of "IoT"
A Q&A with Kevin Ashton

10
The IoT Landscape

12
The Future of Inertial Motion Tracking
By Peter Li

16
Unlocking the Value of the Internet of Things
by Peter Saint-Andre

21
The Promise of Predictive Analytics
by Brad Keywell

25
Unpowering the Internet of Things
by Joshua Smith

29
From the IQT Portfolio

ON OUR RADAR

Making IRL* Computable:

The Inevitable Internet of Things

by Ravi Pappu

* In Real Life

The Internet of Things (IoT) can be defined as the collection of physical objects that communicate their identity, state, and location to the internet. Much ink has been spilled about the number of things this collection might contain in 2020[1], the economic impact that the growth of IoT portends, and the benefit that IoT could have for our industries, cities, farms, vehicles, and even our bodies. However, those discussions offer little insight into how to think broadly about IoT systems, how these systems are put together, or how they might evolve. We consider these questions in this article. In service of clarity, we have omitted specific types of IoT sensors in this article, as that discussion does not necessarily enhance understanding of the bigger ideas at play.

What is IoT?

If you were to X-ray any of the numerous IoT systems in use today, you would see a skeleton that looks like Figure 1. Systems are comprised of *nodes*, which live in the physical world; *gateways* that enable the aggregation and forwarding of data from nodes; and the *internet*, where data is stored, analyzed, and made available for further consumption. This simple description belies the enormous complexity of these systems, but suffices to illuminate the dominant paradigm.

Before we explore the key ideas and the deeper structure of IoT systems, we present a quick look at IoT through the lens of prominent market verti-

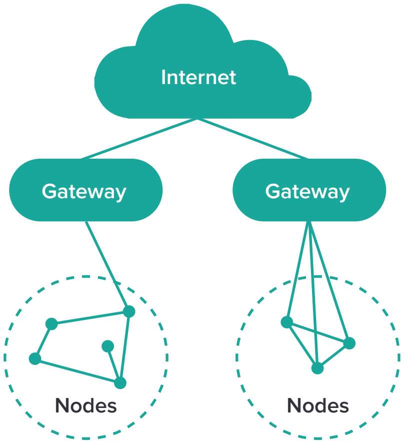

Figure 1 | The dominant paradigm of modern IoT systems, which comprise of nodes, gateways, and the internet.

cals. Our view is primarily informed by examples of what one can do with data collected from IoT nodes.

Market vertical	What is enabled or sensed?	What can be inferred?
Wearables	Physiological parameters, activity duration, location	State of health, psychographic information, patterns of life
Home	Presence, consumption of electricity, water, heat, presence of smoke, fire, carbon monoxide, particulate matter	Home activity, anomalies in consumption patterns, inventory levels
Telematics	Location, speed, bearing, vehicle parameters (fuel consumption, odometer reading, etc.), collisions	Wear and tear on vehicles, adherence to regulations, vehicle maintenance, anomalies in driving behavior, duration of and participation in congregative activities
Commerce	Items purchased and their location	Customer preferences, behavior, connections to payment accounts
Industrial IoT	Physical, chemical, and environmental parameters	Identity of objects, use of transportation systems, supply chain management, efficient use of resources, weather modeling
Robotics	Remote sensing and actuation, driverless cars	Elements of virtual presence, imagery, manufacturing statistics
Smart Cities	Traffic, utilities, waste management, fertilizer	Improved efficiencies, conservation, congestion control
Telemedicine	Physiological parameters, remote health, medication consumption	Population health and wellness parameters, disease prevalence and spread

As is evident from the table, running analytics on sensor data from IoT systems significantly broadens the range of inferences, and thereby, applications. Case in point: a $99 On-Board Diagnostics (OBD) peripheral for automobiles that contains a GPS sensor and can monitor vehicle Controller Area Network (CAN) bus data has enabled dozens of applications ranging from enhancing fuel efficiency to driver safety to expense reporting. It is important to note that the data (and metadata) were always there; the missing ingredients were the low-cost sensor and a means to communicate its data to the cloud. This is obvious, but worth remembering: *IoT makes invisible data visible.*

Fundamental Ideas

The current IoT revolution owes both its genesis and success to four fundamental ideas.

1) Software representations of physical things: Anything that can be represented by software will be represented by software. The first wave of software eating space and time was achieved by thinking deeply about separating the logical content of objects from their physical representation – cleaving the bits from the atoms. With bits in hand, atoms are dispensed with entirely, leading to a slow decline of things like *paper* books, *celluloid* film, and *metallic* coins. Today, even higher functions like operating a vehicle are representable as a combination of sensors and software. This drive toward higher levels of abstraction will no doubt continue relentlessly.

2) Invisible technology: As Marc Weiser so eloquently said, "The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it[2]." Ubiquitous computing — the ability of computing, sensing, and communication technology to disappear into every object and enable those objects to be sensed by computers — has been pivotal in driving IoT forward.

3) Measurement, measurement, measurement: Software is eating the world, but the world can't eat software. Some of the largest challenges facing our species cannot be solved by code; we cannot program away climate change, water contamination, crowded cities, or hunger. However, measuring relevant quantities of interest can help model and understand these complex problems. To this end, IoT offers a scalable, effective, inexpensive, and persistent way of measuring a vast range of quantities.

4) Recombinant[3] technology capabilities: The more technology artifacts we have, the more we *will* have, owing to the power of recombination. Engineering relies on the encapsulation of discrete capabilities into modular artifacts that can then be combined to create new artifacts, which can themselves be modularized, and so on. Modern IoT is primarily based on synthesizing new applications by assembling existing technology capabilities in new ways and uniquely challenging their limits.

These four ideas make IoT inevitable. If we had not yet invented the idea of IoT, we would have to do so now. At this point, we are also in a position to define IoT succinctly: *the Internet of Things makes the real world amenable to computation.*

Six Core Capabilities of IoT Systems

IoT systems require six core technology capabilities, each of which is uniquely challenged by IoT applications. Innovation in any of these capabilities has the potential to significantly broaden the reach of IoT.

1) Communications: IoT needs radios that enable long-range communications at low data rates. It would not be an exaggeration to say that advances in low-cost and low-power radio communication have been pivotal in accelerating IoT deployments. While earlier generations of IoT systems were tethered to the internet with wires, the wireless communication revolution has had a direct causal impact on enabling, to paraphrase a cell phone commercial[4], more IoT in more places. The key requirements for IoT applications are:

- Long-range (miles, not feet) at low power consumption (milliwatts)
- Protocols optimized for short data payloads at low duty cycles as opposed to, for example, 4G-LTE, which is optimized for high bandwidths being utilized continually
- Support for mobility and in-building penetration
- A very low cost approaching $1 per module

In Figure 2 (opposite page), we take a look at the landscape for existing communication protocols as well as emerging protocols that are dedicated to IoT. As is clear from the figure, there are many incumbents[5] for short-range communication, but there is plenty of opportunity for long-range, low-bandwidth systems. This is the space that efforts like SigFox, Ingenu, LoRa, and the other Low-Power Wide Area Networking (LPWAN) efforts are targeting. However, cellular incumbents are developing flavors of LTE to bridge the gap from 4G to 5G, which promises support for IoT requirements from the get-go.

2) Hardware: IoT needs hardware that is low cost, low power, interoperable with a wide variety of sensors, and packaged in rapid prototyping kits to enable quick-turn application development.

Hardware is required to do several different things: general purpose computing, analog-to-digital conversion, storage, and communication. Node hardware spans many orders of magnitude in clock speed, cost, and capability – from a 7-cent passive RFID tag[6] to a portable weather station running on an 8- or 16-bit microcontroller at 20MHz. Node hardware has benefited tremendously from:

- **Moore's Law:** number of transistors per unit area doubles every 18 months
- **Koomey's Law:** number of instructions per joule grows more than Moore
- **Kryder's Law:** amount of storage per unit area also grows more than Moore

The convergence of scaling laws and the maker movement[7] has given rise to a large ecosystem of hardware for IoT applications in the form of single board computers such as Arduino, Raspberry Pi, Beaglebone, and Gumstick.

3) Software: IoT needs modular software that will run on different hardware architectures in resource-constrained environments (with low memory and clock speed). This

IQT Technology Architectures

A Technology Architecture is a unified, coherent structure that shows constituent technology capabilities and how those capabilities fit and work together.

In-Q-Tel's Technology Architecture Group is implementing a model on which Technology Architectures will be developed and used as a foundation for strategic investing against our customers' mission priorities. IQT's model for creating architectures is inspired by object-oriented design, which extends the principles of abstraction and re-use to define core technology capabilities.

By identifying capabilities and decomposing a system into individual core technology components, IQT is able to have more meaningful dialogue with customers. Technology conversations become easier and more contextual - we can discuss smaller components of the problem while keeping the holistic context of the problem intact. Legacy customer solutions and technologies can be mapped to the architecture in order to categorize and compare with potential solutions that exist in the commercial market.

software must be supported by a comprehensive menu of APIs, libraries, and wired/wireless networking stacks.

We are seeing these challenges addressed in two distinct ways: *top-down* and *bottom-up*. The top-down approach involves paring down existing operating systems like Linux and Android to fit the needs of IoT applications. For instance, Google's Brillo[7] is Android stripped down to run on constrained processors. However, these top-down efforts don't quite meet necessary resource constraints. The alternative bottom-up approach is essentially building an IoT operating system from scratch. There are a large number of such OSes available, the most prominent of which are Contiki, TinyOS, and Riot.

In addition, there are several industry alliances coalescing around different IoT verticals that aim to standardize interoperability between IoT devices. These projects, e.g., IoT@Work, Alljoyn, IPSO, and the Open Interconnect Consortium, are usually sponsored and led by large companies with current or future products in those verticals and are still in early stages of development.

4) Security: IoT needs security but it has often, and with very good reason, been called the Internet of Insecure Things[8] with "hilariously broken"[9] security. Nothing exemplifies this characterization more than a cursory browse through Shodan or Censys[10], search engines for IoT devices, which reveal vast numbers of devices without security controls. Some of these devices are innocuous, but many unsecured devices violate privacy or have the potential to cause severe disruption[11].

There are several reasons for this state of affairs. Building secure systems is challenging in any situation, and is exacerbated by the fact that IoT systems run in resource-constrained environments and are frequently deployed in remote locations by non-security professionals. The attack surface is simply too large.

In general, IoT security can be approached in two ways. The first is to optimize existing, well-understood cryptographic standards for operation on IoT nodes; the second, christened Lightweight Cryptography (LWC)[12], is to develop new cryptosystems for such devices. The security prescription for IoT devices can be stated quite simply: Use existing NIST standards wherever possible, because LWC is still in its infancy. We summarize the security landscape for IoT in Figure 3.

Fortunately, all of the major chipset vendors and IoT operating systems support NIST standards in their standard offerings. Some vendors are also offering trusted computing platforms like the Trusted Platform Module and Trusted Execution Environment in their products.

5) Management: With IoT, device management assumes outsized importance. There are three main reasons for this:

1) There are a multitude of nodes which are deployed in remote locations and only occasionally connected to the network;

2) The range and possibility of errors at the node are large; and

3) The cost of mismanagement (e.g., a botched over-the-air upgrade) is a "bricked"[13] node.

IoT device management can range from simple data collection to preemptive failure prediction with typical functions including initial device provisioning, firmware management, device monitoring, logging, and con-

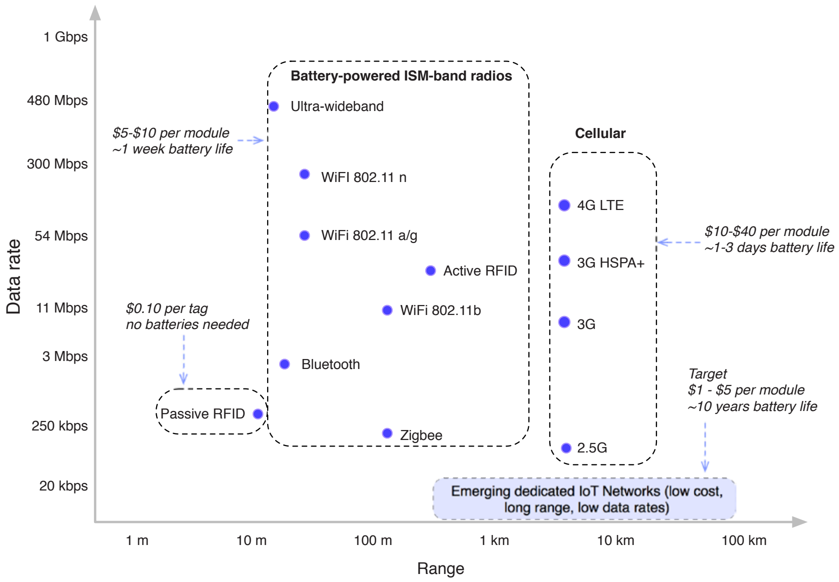

Figure 2 | This chart shows how IoT radios are dominated by short-range, high-bandwidth systems, while there is a need for long-range, low-bandwidth, and low-cost systems for many applications. These systems are depicted in the blue box in the figure. Note: this is not an exhaustive list of protocols and networks.

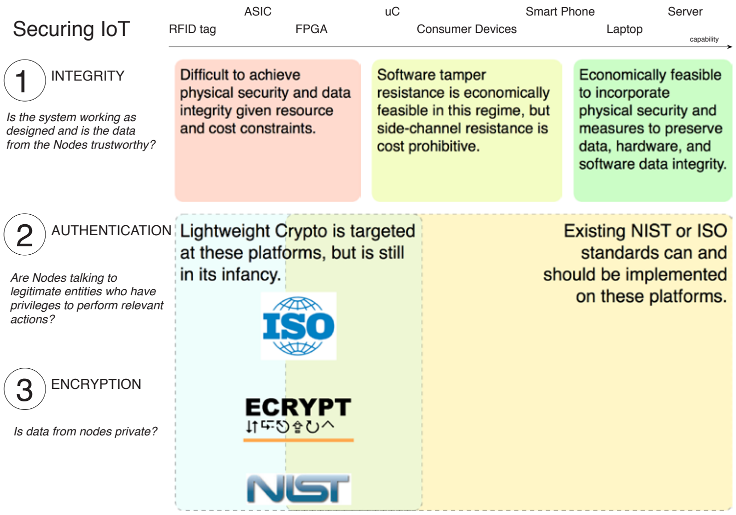
Securing IoT
ASIC
uC
Smart Phone
Server
RFID tag
FPGA
Consumer Devices
Laptop
capability
1 INTEGRITY
Is the system working as designed and is the data from the Nodes trustworthy?
Difficult to achieve physical security and data integrity given resource and cost constraints.
Software tamper resistance is economically feasible in this regime, but side-channel resistance is cost prohibitive.
Economically feasible to incorporate physical security and measures to preserve data, hardware, and software data integrity.
2 AUTHENTICATION
Are Nodes talking to legitimate entities who have privileges to perform relevant actions?
Lightweight Crypto is targeted at these platforms, but is still in its infancy.
Existing NIST or ISO standards can and should be implemented on these platforms.
ISO
3 ENCRYPTION
Is data from nodes private?
ECRYPT
NIST

Figure 3 | The state of IoT security. The axis at the top represents capabilities of IoT hardware. There are three main questions: Can we trust data from IoT nodes (integrity)? Are they talking to legitimate entities (authentication)? Is the data private (encryption)? For really constrained devices, trust is difficult to achieve, whereas for the medium and higher capability devices, it is technically feasible, but might not be economically feasible. With NIST standards, authentication and encryption are feasible in all but the lowest capability devices. Hence we prescribe using these existing standards wherever possible. Light Weight Crypto (LWC) is still in its infancy, but has several major contributors including NSA.

trolling nodes on demand (usually for troubleshooting purposes). We are seeing three approaches to IoT node management:

- Hardware vendors like ARM and Telit are including management capabilities as part of their product line. This includes software on the nodes to connect directly to their cloud platforms for management.
- Cloud computing incumbents like Amazon Web Services, Google Compute Engine, and Microsoft Azure are providing hardware-agnostic methods of getting IoT data into their cloud platforms.
- Hardware prototyping platforms like Arduino have management capabilities ported to run on them.

6. Analysis: IoT systems need analytics to aggregate noisy, granular data from the field and turn it into useful insight. This is what customers pay for.

In most IoT systems, analytics is where big data and data science collide. Many big data techniques are having a major impact on IoT data processing, including:

- New data abstractions for streaming data and distributed stream processing frameworks
- High-performance distributed data stores including NoSQL, time-series, in-memory, graphics processing unit (GPU), field-programmable gate array (FPGA), and geospatial databases
- Probabilistic algorithms
- Machine and deep learning
- Domain-specific analytic frameworks for applications including geospatial, transportation, agriculture, and mining

These core capabilities are visualized on page 10.

The Future of IoT

Predicting the future is hard, but it's clear that there is a fortune of economic and societal value at the bottom of the IoT pyramid[14]. As we said previously, some of humanity's hardest problems could leverage IoT to aid in understanding their scope and pointing the path towards solutions. Beyond the economic value and the number of devices, what can we expect to see as this IoT revolution unfolds? We conclude with one possibility.

Given the relentless drive towards software representations of everything, imagine that the cloud contains more and more sophisticated models, avatars, if you will, of all things IoT. Each of these models is occasionally in contact with its physical twin to refresh its state, but most of the commerce and transactions of data between these objects is happening predominantly in the cloud. What if these models include complete representations of farms, factories, vehicles, and cities? How might that change our economies, and, indeed, our world? Q

Dr. Ravi Pappu *is a Principal Architect in In-Q-Tel's Technology Architecture Group. Prior to joining IQT, Pappu held senior technology and management positions at Trimble Navigation and ThingMagic, a venture-backed company he co-founded. ThingMagic was acquired by Trimble in 2010. Pappu received his Ph.D. from the MIT Media Lab in 2001, and was named to Technology Review's TR100 list of top innovators under 35 and Boston Business Journal's 40 Under 40. He is no longer accepting any age-revealing awards.*

References

1. Estimates range from 21 to 50 Billion IoT devices by 2020
2. Weiser, Mark. The Computer for the 21st Century. Scientific American. September 1991.
3. Inspired by Recombinant DNA: DNA molecules formed by bringing together genetic material from multiple sources, creating sequences that would not otherwise be found in the genome.
4. Gardiner, Bryan. AT&T: 'More Bars in More Places' Is the New 'Fewest Dropped Calls'. Wired.com. http://www.wired.com/2007/08/att-more-bars
5. Morris, Iain. Vodafone to 'Crush' LoRa, Sigfox With NB-IoT. LightReading.com. http://www.lightreading.com/iot/vodafone-to-crush-lora-sigfox-with-nb-iot/d/d-id/722882
6. RFID Frequently Asked Questions. RFID Journal. https://www.rfidjournal.com/faq/show?85
7. Brillo. Google Developers. https://developers.google.com/brillo/
8. The Internet of Insecure Things. Forbes.com. http://www.forbes.com/sites/moorinsights/2015/09/.../the-internet-of-insecure-things/
9. Porup, J.M. "Internet of Things" security is hilariously broken and getting worse. Ars Technica. http://arstechnica.com/security/2016/01/how-to-search-the-internet-of-things-for-photos-of-sleeping-babies/
10. Shodan. https://www.shodan.io
11. Snow, John. What are IoT search engines Shodan and Censys and what are they capable of? Kaspersky Lab Official Blog. https://blog.kaspersky.com/shodan-censys/11430/
12. Lightweight Cryptography Workshop 2015. http://www.nist.gov/itl/csd/ct/lwc_workshop2015.cfm
13. Bricked: To render completely useless, as useless as a brick.
14. Inspired by: Prahalad, C.K. and Stuart L. Hart. The Fortune at the Bottom of the Pyramid. http://www.stuartlhart.com/sites/stuartlhart.com/files/Prahalad_Hart_2001_SB.pdf

A Look Inside the Issue

This issue of the *IQT Quarterly* examines the Internet of Things (IoT) revolution.

We begin with a Q&A with Kevin Ashton, a pioneer who is credited with coining the term "Internet of Things." He discusses the circumstances that led to his coining the now widespread term, and his experiences working on a global RFID standardization effort at MIT.

Peter Li discusses how Atlas Wearables is doing for human motion what Siri did for human voice, but without a tether to the cloud. The key ingredients of the Atlas recipe are low-cost, low-power 3D motion sensors and a learning system that runs on commodity microcontroller.

We then switch to the gateway layer of IoT. Peter Saint-Andre of Filament discusses a unique peer-to-peer, long-range radio communication system that uses smart contracts and private microtransactions to communicate and exchange value in a completely decentralized fashion.

Brad Keywell of Uptake then explains how to turn raw data from IoT sensors into knowledge by using modern big data technologies in the analytics layer. This layer is where domain-specific algorithms meet noisy, high-velocity data from the field.

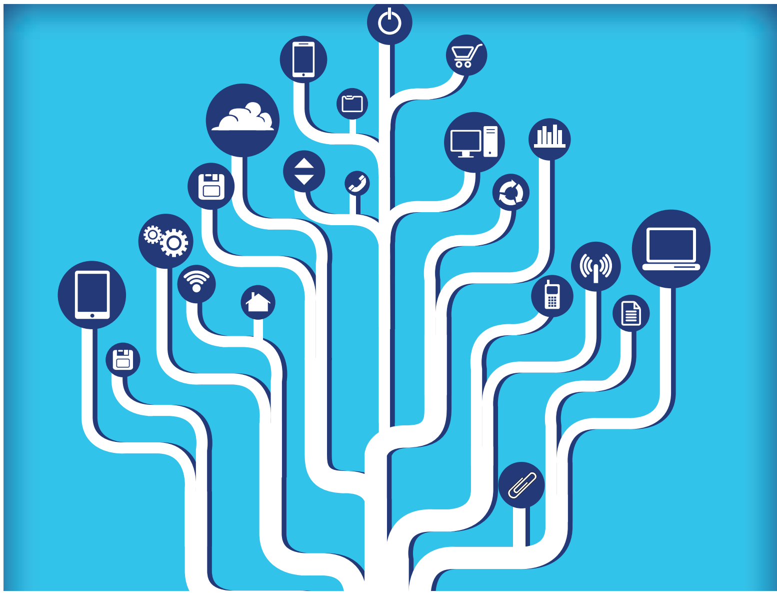

Finally, we look ahead to a crucial challenge that must be solved for IoT to expand its reach even further: energy. Josh Smith of the University of Washington presents the results of a decade of research that use the principles of passive RFID tags and show how they can be leveraged to harvest increasing amounts of energy from radio frequency transmissions. The devices range from an accelerometer to a microphone to a camera, all of which are powered solely by ambient RF. **Q**

The Genesis and Evolution of "IoT"

A Q&A with Kevin Ashton

Kevin Ashton is a technologist credited with coining the term "Internet of Things." An entrepreneur who has led several successful startups, Kevin also is known for his writing on technology and in 2015 published a book, "How to Fly a Horse: The Secret History of Creation, Invention, and Discovery." IQT's Ravi Pappu recently spoke with Kevin about the genesis of the term, the evolution of IoT, and innovation.

What were the circumstances that led you to coin the term "Internet of Things"?

That happened in the late 1990s. I was an Assistant Brand Manager at Procter & Gamble—my first job out of school—and I was part of a team launching a new range of make-up products. We had one particular shade of lipstick that was very popular, partly because we were advertising it, and it was always out of stock at my local store when I went to get my weekly groceries. That p****d me off. So I investigated, and after several months of work, found it was out of stock in 4 out of 10 stores at any given time, and that, furthermore, Procter & Gamble's most advertised products were nearly always its most out of stock products. That was an interesting—and expensive—supply chain problem. We knew advertised products would sell more, so of course we always made extra, but that wasn't solving the out of stock problem.

I did more digging, and eventually realized the problem was information, which was a bit of a surprise, because the 1990s were supposed to be an "information revolution." The fundamental issue was that, in the twentieth century, all digital information was entered by human beings, with only a few exceptions. For people used to the twentieth century paradigm of computing—which is to say, pretty much anyone born before about 1990—that can be a difficult problem to grasp. Many of them respond to that statement fairly blankly: they simply don't get it, and wonder, "What other kind of information is there?" In fact, it's a profound limitation, because there are only a few types of information that human beings are good at entering—say, conceptual data, like payrolls, or appointments, or ideas and thoughts—and billions of types of information that human beings simply cannot enter, such as detailed data, or data that changes constantly, which is to say pretty much all data about the real world. And that includes whether a particular shade of lipstick is on the shelf in a particular grocery store at a particular moment.

So, the big information technology question at the start of the twenty-first century was, "How do you gather data that cannot be gathered by human beings?" Like all good questions, the question answers itself: you build automated systems that gather data, also known as sensors. And sensors work best in networks—see the human nervous system as an example—and in the mid-1990s, we suddenly had a new worldwide network that would be perfect for connecting sensors: the internet. I had to explain all this to busy, non-technical P&G senior executives, using the corporate medium of choice, PowerPoint slides, so, around the spring of 1999, when coming up with a title for my executive PowerPoint presentation, I boiled it all down to three words: the Internet of Things. That was very weird and ungrammatical at the time—I am not sure anybody had used the phrase "the Internet of" something before, and if they had, if was not widely known—but it caught the attention of senior managers who knew that the internet was some kind of big deal in which they were supposed to be interested and invested.

There are many definitions of IoT out there. How would you describe IoT?

Computers gathering information by and for themselves using networked sensors.

You have seen IoT from several perspectives: as a potential adopter (at P&G), as a standards leader (at MIT), as a vendor (at several startups), and as an influencer (through your writing and speaking career). What are a few things that have surprised you the most about how the field has evolved and why?

I was always the wildest-eyed, craziest guy in the room, and, in retrospect, it turns out I wasn't wild-eyed or crazy enough. We can do things today that were unimaginable, or supposed to be completely impossible, at the turn of the century. The proliferation of high-bandwidth, low-cost radio networks is one example. The sophistication of machine vision systems is another. If you traveled back in time to 2000 and described the technology of 2016, people would think you were mad; experts doubly so, because they would be able to explain why the things you were describing were literally impossible. I know, because I got all those reactions, and the things I was describing, such as 5-cent RFID tags, or always-on, dial-up free internet connections everywhere, were laughably tame compared to today's reality. The biggest surprise is not how fast things change, but how fast everyone gets used to how fast things change. You never find anybody who admits they thought today was impossible yesterday. Everybody is an IoT person now. But 15 years ago, my IoT talks were often met with cold silence, or put-downs that were whispered behind the backs of people's hands, or streams of objections. Not at MIT, where ideas are never damned for being too crazy, only for not being crazy enough, but outside, in the "real" world.

> “I boiled it all down to three words: the Internet of Things...it caught the attention of senior managers who knew that the internet was some kind of **big deal** in which they were supposed to be interested and invested.”

In the next 5-10 years, which technology sectors or applications will see the greatest impact from IoT and why?

Self-driving cars will be the next big one. Those will be here far sooner than anybody seems to realize, and the consequences will be profound, starting with preventing 3,300 needless deaths that humans driving cars cause each day.

You recently wrote a book on innovation titled "How To Fly a Horse." What advice do you have for organizations that struggle with balancing tactical goals with the existential need to innovate?

I don't give advice, but I'll say this: those two things are not in opposition; they are cause and effect. Short-term problems are innovation's seeds. For example, the Internet of Things started with a missing lipstick. It's those gritty, real, frontline nits that give rise to the best ideas, not grandiose business plans, or mythical Archimedean moments. The best businesses seek solutions to their problems, not problems for their solutions. **Q**

__Kevin Ashton__ is a visionary technologist. He coined the term "the Internet of Things," co-founded the Auto-ID Center at MIT, and led three successful tech start-ups, including Zensi, which he co-founded and sold to Belkin in 2010. His writing about innovation and technology has appeared in The New York Times, The Atlantic, Politico, and Quartz.

The IoT Landscape

Internet 3.0: Connecting Everything

The internet (short for inter networking) was born of a need to connect computers to each other. This was followed by an explosion of applications, primarily connecting people to organizations and content and led to expansive growth in search, social networking, and commerce. The emerging IoT revolution promises to connect things in our physical world to the internet, dwarfing both the number of computers and the number of people on the internet by many orders of magnitude.

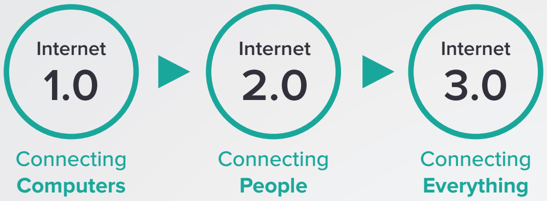

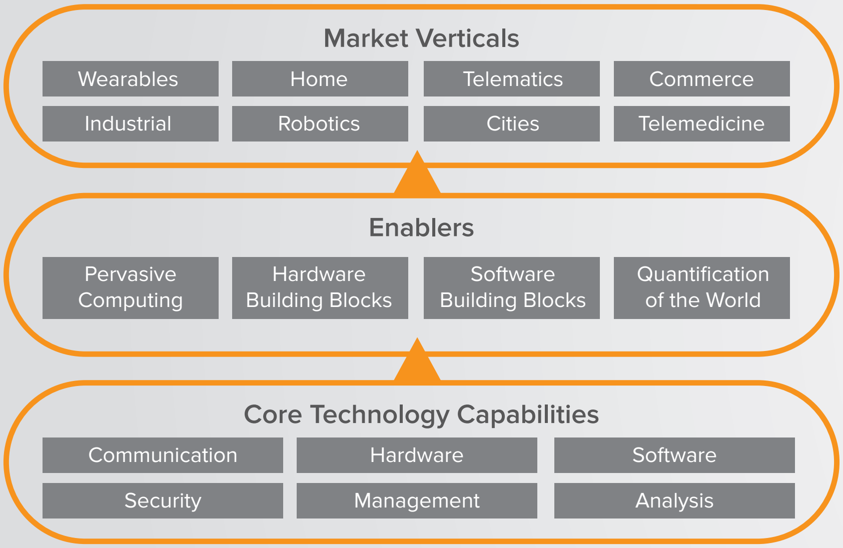

CONNECTED THINGS IN USE WORLDWIDE:

6.4B in 2016[1]

20.8B by 2020[1]

THE INTERNET OF INSECURE THINGS?

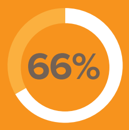

of networks will have an IoT security breach by 2018[2]

A Generic IoT System

The diagram to the right shows the dominant paradigm of modern IoT systems, which comprise of nodes, gateways, and the internet. Nodes can talk to each other and relay messages for other nodes via peer-to-peer communication, to an intermediary called the gateway via short-range communication, or directly to the internet via a cellular or satellite connection. There is a tremendous diversity of nodes owing to the broad range of IoT applications.

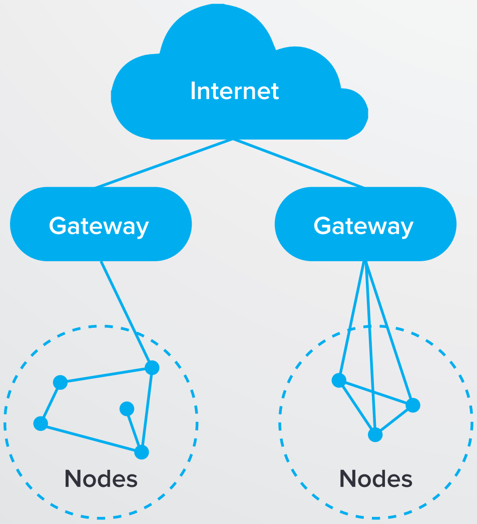

Enabling Technologies	Challenges	Advances
Communications	▪ Low cost ▪ Long range ▪ Low power ▪ Scalability ▪ Interoperability ▪ Efficiency ▪ Mobility	▪ IoT-specific communications protocols ▪ Optimization of existing cellular and short-range protocols
Software	▪ Constrained resources ▪ Scalability ▪ Modularity ▪ Connectivity ▪ Reliability	▪ IoT-specific operating systems (e.g., Contiki, TinyOS) support major multi-chip unit (MCU) families and networking protocols ▪ APIs for everything: sensor integration, device management, message brokering, and more
Hardware	▪ Size and compute power ▪ Battery life ▪ Adaptability ▪ Multi-sensor support	▪ Low-power memory ▪ Near and sub-threshold power ▪ Low-power communications
Security	▪ Usually absent ▪ Budget ▪ Constrained resources ▪ Remoteness ▪ Weak link ▪ Skills	▪ Optimizing NIST standards: reduction in Advanced Encryption Standard (AES) resource requirements, and Data Encryption Standard (DES) modifications ▪ Lightweight cryptography (LWC)
Management	▪ Constrained resources ▪ Scale ▪ Occasionally connected ▪ No downtime ▪ High penalty for failure	▪ Device services: registries, discovery, and search ▪ Purpose-built IoT platforms with APIs and SDKs that allow for node management ▪ Full stack development tools
Analysis	▪ Cost ▪ Range ▪ Power ▪ Scalability ▪ Interoperability ▪ Efficiency ▪ Mobility	▪ New streaming data abstractions ▪ Probabilistic algorithms ▪ High-performance distributed data stores ▪ Commoditization of machine and deep learning

The Future of Inertial Motion Tracking

by Peter Li

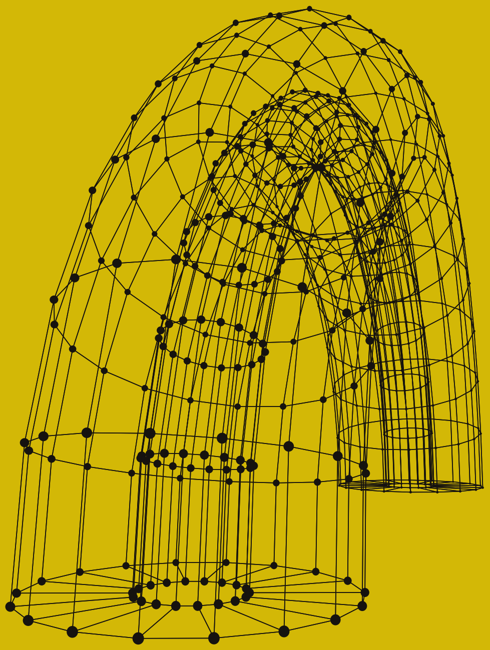

The inertial microelectromechanical sensor (MEMS) market has expanded rapidly in the last half-decade. Following the advent of smartphones, inertial MEMS proliferation has been driven by the demand for hardware design optimizations and breakthroughs in new detection principles. As more and more devices are connected, new opportunities for increasing functionality and end user value emerge.

So far MEMS technology has been limited to general activity and location functionality. This is useful for the most basic tasks as a new mode for user interface design or basic activity detection, but not much more. Consider a deployment in a physical rehabilitation program. Prescription, compliance, and improvement are key variables inherently linked to patient outcomes. While troves of inertial MEMS data can be gathered, unearthing hidden value is difficult. To make this possible, we need a technology stack that digests this data to enable automation of tedious tracking and deeper analysis with these data points. Elements of the rehabilitation program, including progress, range-of-motion, and stability, can be measured and calculated automatically. In this article, we show how these interactions and elements for analysis are enabled by off-the-shelf hardware and software that emphasize computational and power efficiency and untethered functionality.

Background

Inertial navigation was first pioneered by navigation of aircraft, tactical and strategic missiles, spacecraft, and other military applications. Recent advances in the manufacturing and deeper understanding of physics have made it possible to manufacture small and light microelectromechanical systems. These innovations have broadened the spectrum of possible applications. We first experienced this at full-scale in 1993, when accelerometer MEMS were first used in consumer vehicles to trigger airbags. In the last decade, MEMS components experienced explosive volume growth in parallel with increasing power efficiency and decreasing costs due, in part, to the introduction of the modern smartphone in 2007.

This commoditization has helped enable innovations in wearable devices and the Internet of Things (IoT). Today, inertial MEMS are designed into products including watches, glasses, shoes, and tools. Many technologies today leverage this new, affordable source of data to categorize five general modes of motion sensing: acceleration, vibration, shock, tilt, and rotation. Although these modes are valuable as new inputs for user interface and experience design, they are inherently limited by the inability to analyze the raw data. Since this data is not detailed or sufficiently refined, analysis and knowledge extraction is hindered significantly – garbage in, garbage out. Atlas Motion Engine technology alchemizes this garbage data from MEMS devices by identifying and analyzing key patterns in 3D motion.

Motion Engine Technology

Atlas Motion Engine technology is optimized for deployment on low-power microcontroller technology connected to 6-axis inertial MEMS. Today, Atlas comes pre-loaded with over 70 fitness-related activity models and adapts to users and even learns new motions through cloud training.

Keeping resource constraints on the wearable device in mind, the Atlas platform is built on three key principles:

1. **Power efficiency:** Power efficiency and power consumption are perhaps the most important driving factors when developing connected devices that require a portable source of power. Many deployments prefer and prioritize solutions that require less charging and can run longer on a single charge.
2. **Computational efficiency:** Pseudo real-time detection and a small software footprint are key requirements in many use cases. For example, real-time identification of falls by elderly people on a small microcontroller is ideal.
3. **Untethered / cell service denied:** The ability for untethered detection is driven by many use cases where instant results are a necessity or cellular service is unavailable. Most voice recognition solutions require a constant connection to a cellular signal to outsource computations to a remote server. This tethering not only requires significantly higher power consumption but also introduces unwanted latency. Unlike many other machine learning or classification platforms, Atlas Motion Engine is fully embedded on a microcontroller.

Advanced motion tracking consists of two main components: the embedded Engine and cloud-based Engine. The embedded Engine is comprised of motion classification capabilities and analytics. The cloud-based Engine provides additional and more fine-grained metrics as well as enables learning and training of new custom 3D motions. Once the new 3D motions are learned, they can be transferred to the embedded Engine for an untethered tracking experience.

We take a hybrid approach to learning activity models. Cloud-based learning takes advantage of the benefits of discriminated inertial data and the large amounts of available compute power in the cloud to generate the most accurate activity models. The learning protocol is generally batch-processed to “imprint” the chip with the user's preferences for tracking. Once updated, a typical deployment does not include constant adaptation.

New Elements Enabled by Atlas Motion Technology

The primary advantage of motion-based activity classification is that it is a simple, efficient, and completely wearable means of evaluating 3D motion. Motion-based activity classification is proven to monitor over a hundred distinct motions for long periods and can be as simple as a single wearable device. While hundreds of full body motions can be effectively classified and analyzed from a single point sensor, multiple sensors can further improve the accuracy and breadth of classification.

Motion-based activity classification also creates an incredible source of data on human health responses to fitness regimens. Various studies[1,2,3] conclude that this resolution of data is more than enough to obtain fit-

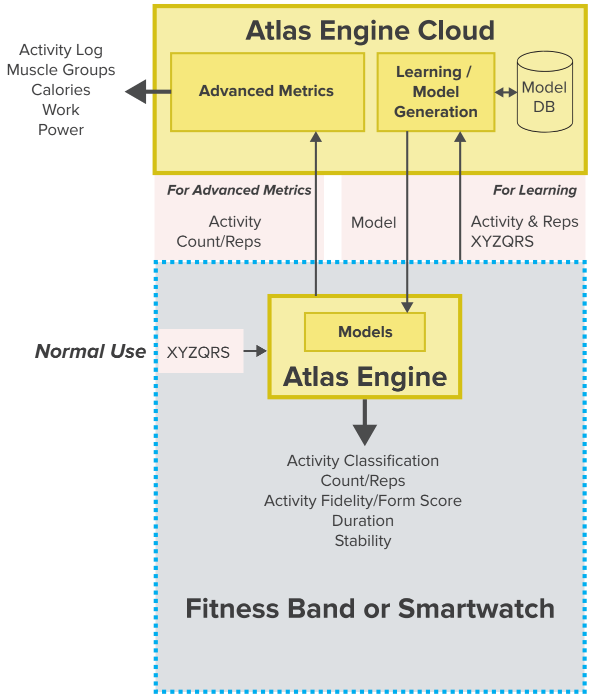

Figure 1 | When considering inertial motion sensors with six degrees of freedom, systems can measure cartesian acceleration along three axes (X, Y, and Z) as well as rotational angular acceleration along three axes (Q, R and S).

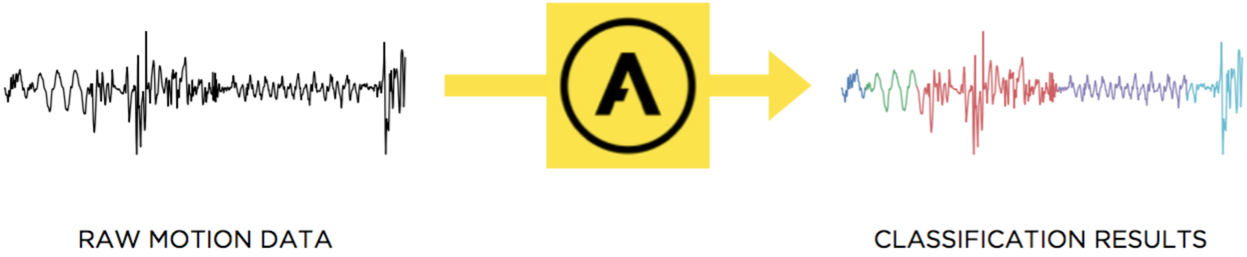

ness progress and specific metrics that can help predict physiological response. Atlas technology provides many metrics such as stability, motion fidelity, and velocity. Applications range from personal fitness to workforce or insurance compliance, among many others.

The technology can expand to support multiple specific use cases. In the fitness case, the technology enables automatic calculation of metrics like reps, rest time, and velocity. Below, we take a look at how inertial sensing can be used to measure minute details of workout regimens in order to improve their efficacy.

Physical performance and physiological adaptations are demonstrated to be linked to the intensity and number of repetitions performed[1]. In general, studies divide subjects into four groups of repetition training; a low repetition group (3-5 reps), an intermediate repetition group (9-11 reps), a high repetition group (20-28 reps), and a control group. Low and intermediate rep groups are shown to have significant hypertrophy for all major fiber types (types I, IIA, and IIB). The high rep group is better adapted for submaximal, prolonged contractions that improve aerobic power and time to exhaustion.

Rest interval between sets is an important variable in resistance exercise prescription. The amount of rest between sets can influence the efficiency, safety, and ultimate effectiveness of a strength training program. Higher levels of muscular power have been demonstrated across sets with 3-5 minutes of rest versus 1 minute of rest between sets[2]. These findings indirectly demonstrate gains in muscular endurance when utilizing short rest intervals.

Velocity has been used as a fitness metric for a few decades, but has only recently grown in popularity. The researchers who discovered the concept were trying to understand what optimal weight should be used for a variety of training exercises, and they used velocity of the barbell to determine the weight parameters of the load. For athletes competing in sports, the use of velocity-based resistance training has been shown to improve performance on sport-specific tests and much of that has to do with the muscle types involved, not just in the sport, but also the position. For example, fast lengthening contractions are attributed to greater hypertrophy and strength gains compared to slow velocities[3]. Type I muscle fiber size increased in both fast and slow training. However, type IIA and IIX muscle fiber cross-sectional area increased in both types of training but the increases were greater with fast velocity training. Different muscle composition adaptations are pivotal to reaching specific and personal fitness goals.

Application Exploration

The Atlas platform has been extended to support a wide variety of motions. For example, when sensors are incorporated into a toothbrush, or the dominant hand, one can track and identify brushing coverage. Top, bottom, left, and right mouth tracking can be discerned to help users understand the efficacy of their oral hygiene.

Consider a deployment in a shipping and fulfillment organization, where line workers lift heavy boxes. The number of lifts can be tracked automatically. But beyond simple counting and tracking, the Atlas metric solution can be adapted to identify fatigue over time, improper lifting form, and other key metrics to help managers quantify efficiency and to help line workers identify potentially risky patterns.

Conclusion

The Atlas Motion Engine does for human motion what natural language processing (NLP) did for human voice. In NLP, a crucial ingredient required to accelerate the field was the digital microphone (sensor) and the algorithms (analytics). Without the algorithms, all we had

> “The growth of voice processing and NLP enabled a brand new generation of technology controlled by voice. This is how we expect human motion tracking to evolve.”

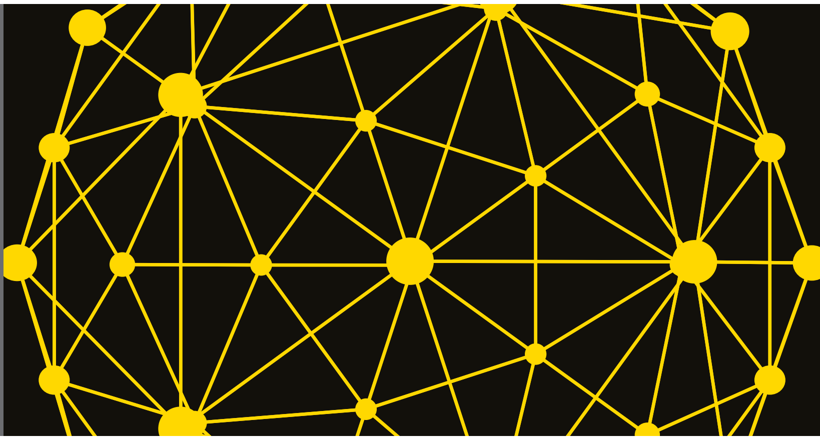

was a mess of raw data. The growth of voice processing and NLP enabled a brand new generation of technology controlled by voice. This is how we expect human motion tracking to evolve.

It remains to be seen whether an organization can effectively leverage raw inertial motion data beyond simple pedometer tracking. Atlas is empowering applications to create novel value for their users. This article has provided a high-level view of implementation organization and the technologies required. Although there is much work to be done to broaden the inertial motion database, many important motions and patterns are already included. Looking ahead, additional applications where automatic motion classification and automation can add new interpretations of raw data and enable a new generation of products. Q

Peter Li *is a biomedical engineer from Johns Hopkins University and co-founded Atlas Wearables. His background and focus revolve around data analytics and machine learning techniques. He swam for over 14 years of his life and believes in the automation of tedious chores by way of advanced software solutions.*

References

1. Campos, G.E.R. et al. Muscular adaptations in response to three different resistance-training regimens: specificity of repetition maximum training zones. Department of Biomedical Sciences, College of Osteopathic Medicine, Ohio University. Eur J Appl Physiol (2002) 88: 50-60.
2. Anderson, L. et al. Changes in the human muscle force-velocity relationship in response to resistance training and subsequent detraining. Institute of Sports Medicine Copenhagen/Team Danmark Testcenter and Copenhagen Muscle Research Center. J Appl Physiol 99; 87-84, 2005.
3. Shepstone, T. Short-term high- vs. low-velocity isokinetic lengthening training results in greater hypertrophy of the elbow flexors in young men. Exercise and Metabolism Research Group, Department of Kinesiology, McMaster University. J Appl Physiol 98: 1768-1776, 2005.

Unlocking the Value of the Internet of Things

by Peter Saint-Andre

As more and more devices are connected in the Internet of Things (IoT), an enormous amount of value is waiting to be unlocked. The imminent prospect of extending connectivity to trillions of devices opens the possibility of what we might call picoeconomics, a thousandfold increase in value exchanged over the nanoeconomics implicit in the interactions between billions of human beings on our planet.

To help make this new world possible, Filament has built an open technology stack called Distributed Sentient Transactions (DIST). DIST leverages the most advanced communication and security methods available today and thus enables devices to discover, communicate, and interact with each other in a fully autonomous and distributed manner.

Consider a deployment on a rail network, in which locomotives, freight cars, switch motors, and other pieces of infrastructure are networked through inexpensive, surface-mount devices or onboard firmware. Suitably designed, such devices can communicate with each other over long-range radio at distances up to 10 miles instead of relying exclusively on WiFi, cellular, or satellite access. Now the rail network can gather real-time data from all of these devices under a variety of network conditions, run preventive analytics in the cloud, change the behavior of edge devices by deploying updated directives, perform targeted maintenance, and reduce the risk of dangerous and costly accidents.

Beyond mere connectivity, the devices involved can exchange value directly or indirectly with a wide range of entities. For example, they could sell data about environmental conditions to a meteorological agency, usage data of the rail network to an organization specializing in business statistics, monetize access to their private communication network to customers along rail routes such as grain elevators and loading docks, or productize access to the rail cars directly. As these examples indicate, the exchange of value is not limited to a single location or vertical but can cross organizational lines in secure and flexible ways.

Principles of IoT Development

At Filament, we believe that applications such as these must be built upon five key principles: Security, Privacy, Autonomy, Decentralization, and Exchange (SPADE).

The security principle guarantees that information is not disclosed to unauthorized entities (confidentiality) and not modified in an unauthorized or accidental manner (integrity). Security often involves encryption: encoding

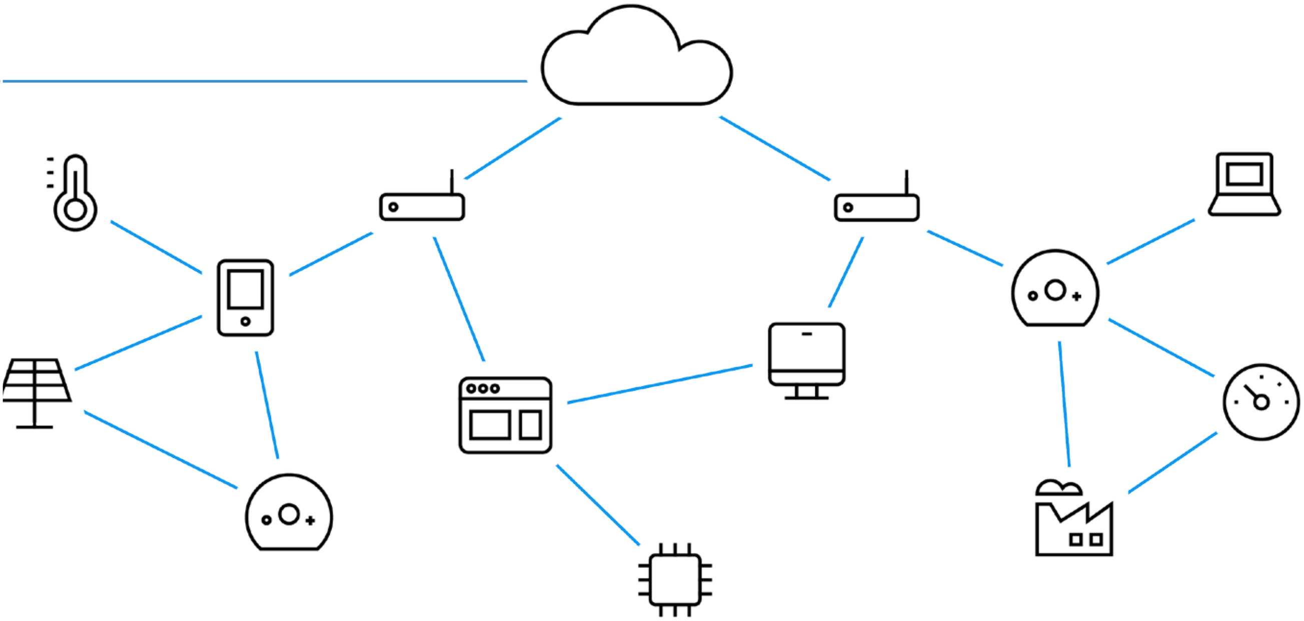

information so that only authorized entities can decode and thus understand it. IoT technologies need to natively ensure the confidentiality and integrity of information, not depend on aftermarket workarounds.

Privacy involves the protection of information about interactions, as opposed to the interactions themselves. Such information (often called metadata) might enable an attacker to correlate interactions with a particular individual, analyze the traffic generated by an endpoint, uniquely identify or "fingerprint" a device, or otherwise detect the identity or attributes of an entity. Privacy-respecting technologies prevent attackers from learning such information, for example, by using ephemeral addresses or routing data over ad-hoc links.

Autonomy means that devices are "first-class" citizens on the internet and are not dependent upon centralized servers. For example, the cars of a freight train should not need to call out to a cloud service in order to communicate with each other or with the head-of-train device in the locomotive. Instead, devices need to be empowered to interact whenever necessary, even if the cloud is unavailable.

Decentralization is a direct result of device autonomy, because no central authority regulates or makes decisions for the actors on a given network. Popular internet services are monocentric: endpoints and their data are tied to each service, and the services do not talk to each other. Some existing technologies (e.g., email) are polycentric: anyone can run their own service on a federated network where services talk to each other, but the services are still primary and endpoints are secondary. By contrast, autonomous endpoints can self-form their own networks, thus building up completely decentralized architectures of communication. This is important for devices in remote locations, but it is also the key to unlocking value from device-to-device interactions.

The foregoing building blocks enable devices to interact in completely independent ways. One further ingredient

The SPADE Framework

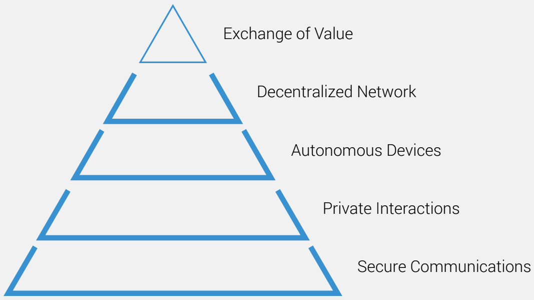

– the capability to engage in smart contracts – makes it possible for such devices to leverage all of the societal and legal constructs for economic exchange that have been built up over thousands of years (such as binding agreements, bills of sale, and validated receipts). The values exchanged through smart contracts and microtransactions among IoT devices could include data, network access, currencies such as Bitcoin, compute cycles, contracts for ongoing service, trusted introductions to other devices, and much more.

From Principles to Features

The SPADE design criteria animate all of the core features of the DIST stack: communicating with other devices, discovering their identities and capabilities, negotiating interactions, and finally, exchanging value.

DIST uses a protocol called telehash for communication among devices. Unlike technologies that rely on centralized or federated servers, telehash enables completely distributed, decentralized interaction. To ensure confidentiality and integrity, telehash messages are always 100 percent end-to-end encrypted using forward secrecy and advanced stream ciphers. For privacy protection, built-in cloaking mechanisms can add random noise to all bytes sent across the wire. Although telehash is transport-agnostic and can be run over standard protocols such as TCP and HTTP, the TMesh extension to telehash provides packetization of telehash over sub-GHz, long-range radio (including shared management of available spectrum and establishment of networking relationships among deployed devices).

Anatomy of a Smart Contract

HEADER	Hashing Algorithm	
PAYLOAD	RESERVED	Issuer, Expiration, Audience, etc.
	PUBLIC	Standard Values (e.g. OpenID Connect)
	PRIVATE	Filament-specific Values
SIGNATURE	Hash of HEADER + PAYLOAD	

In DIST, identity is tied to an endpoint's telehash address or hashname: a 32-byte string independently generated over a public key, not issued by a central authority such as the Domain Name System (DNS). But how can one device discover other devices and learn more about their capabilities? On a small scale such as equipment on a factory floor, discovery of other devices can occur over WiFi, radio, or even Bluetooth; this kind of discovery happens organically as devices are provisioned into a private community. On a wider scale, the groundwork is being laid for endpoints to be vouched for by recognized notaries.

Discovery of an address is necessary but not sufficient for enabling contracts and microtransactions, because the capabilities of a device will largely shape the range of possible interactions. Such capabilities might include the data types a device can provide, the actions a device can take (such as provide network access), and the affiliations that a device has with other devices, larger communities, and trust sources. In order to protect the discovery of capabilities, the DIST stack uses smart contracts (self-executing, self-enforcing contracts that are implemented in software) via a technology called blocklet. These smart contracts make it possible to specify the particular conditions under which a device will interact with other entities, without reference to a cloud service. Such conditions can include price, the time period during which access is allowed, a per-use charge for defined functionality, attribution for data provided, and other contractual terms that are important to the parties involved.

To enable exchange on top of these foundations, DIST includes a method for secure, private microtransactions among autonomous devices. Blocklet microtransactions, which are based on the IETF's JSON Web Token (JWT) standard (RFC 7519), solve this problem in two ways. First, currency can be exchanged in private side chains that are separate from the public Bitcoin block chain. Second, one or both parties to a microtransaction can agree to use an escrow arrangement as a way to lock the value to be exchanged, such that it can be unlocked only after both entities have fulfilled the terms of the contract.

Although the entities involved might want to engage the services of a third-party auditor for high-value transactions, such an arrangement is completely voluntary and subject to negotiation. In addition, because unlocking

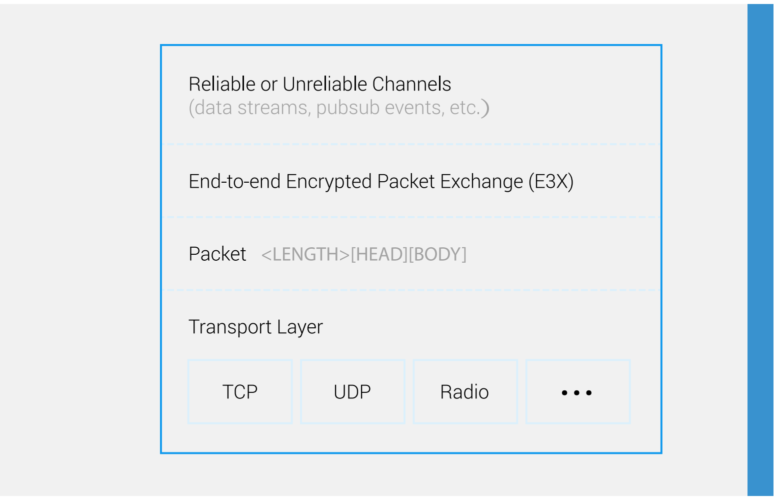

Reliable or Unreliable Channels
(data streams, pubsub events, etc.)

End-to-end Encrypted Packet Exchange (E3X)

Packet <LENGTH>[HEAD][BODY]

Transport Layer

TCP | UDP | Radio | •••

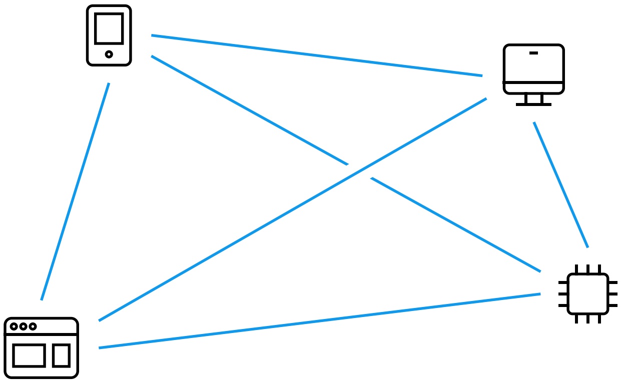

the escrow depends on meeting specified contractual conditions instead of time-based leases, value can be exchanged even if internet connectivity is unavailable for long periods of time.

A New World of IoT Applications

The Filament team has developed DIST not for the technology itself but for practical applications in industrial scenarios, such as energy infrastructure, airports, hospitals, and factories. In many of these situations, privacy and security are not just pleasant-sounding buzzwords but mission-critical necessities. We have also designed DIST so that it can be used where connectivity is intermittent or simply unavailable; pipelines, power grids, oil and gas fields, and mines are just a few examples.

Returning to our example of rail transportation, a number of potential applications arise. For instance, manufacturers such as Hitachi have already started to demonstrate the viability of a "train-as-a-service" model[1], in which the large capital expenditure costs of rolling stock and rail infrastructure are converted into more digestible operating expenses for the customer, who pays only for on-time service. Under this model, it behooves the provider to gather as much information as possible to increase uptime, safely improve delivery times, reduce the risk of accidents, etc. Although end-of-train detectors along with sensors such as hotbox and dragging equipment detectors are standard on rail networks today, inexpensive IoT devices introduce the possibility of monitoring a wider range of data, such as the movement and vibration of individual rail cars (which might indicate, for example, the presence of broken rails or the failure of vehicle running gear).

Naturally, much of the data gathered in this manner will be centrally analyzed using big data methods such as predictive analytics. Yet the gathering itself can be completely distributed across the rail network, with both rolling stock and infrastructure components playing a part. Furthermore, each piece of rolling stock can communicate with others and with the head-of-train device in a locomotive to handle potentially dangerous conditions in close to real time even if a train is far away from conventional internet connectivity. Finally, including a lightweight scripting environment on the end devices introduces the possibility of true edge computing whereby deployed devices can dynamically modify their behavior based on conditions in the field.

With devices deployed throughout the network, truly autonomous markets for information and contractual interaction can be created as well. For instance, an intermodal shipping container might be packed at a factory in China, transported via rail to the port of Hong Kong, transferred to a container ship that crosses the Pacific Ocean, unloaded in Long Beach for transport by rail again, and finally delivered to another factory in Dallas for unpacking there. On each leg of the journey, a device associated with the container can engage in secure microtransactions: selling data about its location and cargo, negotiating connectivity to communicate with its owner, signing off on delivery to its final destination, and perhaps someday even clearing customs.

Many IoT providers wish to provide vertically integrated solutions, from devices on the edge all the way to analytics engines and business decision-making tools. At Filament, we are forging an alternative path: an entirely

> “The exchange of value is not limited to a single location or vertical, but can cross organizational lines in secure and flexible ways.”

decentralized network in which autonomous endpoints use smart contracts and private microtransactions to interact and exchange value in completely voluntary and secure ways. Furthermore, it is Filament's intention to standardize the underlying DIST protocols and provide open source implementations in order to seed innovation and provide a completely open, modular approach to unlocking the value of the Internet of Things. Q

__Peter Saint-Andre__ is a well-known expert on messaging, presence, distributed systems, real-time collaboration, internationalization, and information security. In addition to 20 years of hands-on experience with internet technologies at companies such as Jabber and Cisco, he has been actively involved with industry standardization and has served as an Area Director at the Internet Engineering Task Force. As VP of Strategy at Filament, Peter ensures that Filament's technologies align with customer requirements in a wide variety of industrial IoT applications such as asset management, remote monitoring, and factory automation.

References

1. Yoshida, H. Internet of Things and Train as a Service. Hitachi Data Systems Community. https://community.hds.com/community/innovation-center/hus-place/blog/2014/11/07/internet-of-things-and-train-as-a-service.

The Promise of Predictive Analytics

by Brad Keywell

We are in the midst of a data renaissance. Every single asset and person in an organization is now a data generator – through sensors, computers, smartphones, and even their actions in the real world. Unfortunately, rather than harnessing the unprecedented economic value of this new data, most enterprises are drowning in the data deluge.

Within the next four years, 50 billion machines will be connected to the internet, up from just 12 billion in 2011[1], while data volume is expected to double in size every two years into the next decade as the physical world continues to go online. These conditions are unlocking unprecedented new business opportunities.[2]

To ensure their fair shot at these value-generating prospects, enterprises must solve their data conundrums before it is too late. But with so much data to contend with, how do enterprises know what information is attention-warranting, let alone, actionable? Furthermore, when poor or incomplete data can cost businesses up to 30 percent of revenue each year[3], how are they to ensure the information they've zeroed in on comes from the most reliable sources?

Uptake is doing just that – helping organizations recognize what their most useful data is telling them and effectively predicting the future by transforming messy, unstructured data into insights that compel quick, intelligent actions that increase productivity and enterprise efficiency while enhancing safety and regulatory compliance. We believe that predictive analytics is becoming a critical tool with far-reaching implications for an organization's optimization, or even for its survival.

Rising to Meet Demand in Rail

Nowhere are these implications clearer than in the rail industry. With demand for rail freight transportation on the rise, the industry is faced with significant challenges to better utilize assets and infrastructure to meet these growing demands in an increasingly regulated arena. The greatest of all these challenges is unplanned downtime: locomotive malfunctions, especially those that halt rail lines, cost railroad operators millions of dollars an hour in lost revenue and penalties.

Picture this scenario: Somewhere in the middle of a remote desert, the oil pressure begins to fall on a locomotive hauling millions of dollars of freight. Left unaddressed, this problem could quickly lead to an engine failure, stranding the locomotive on a remote stretch of track, delaying other critical freight missions online until the failed asset can be hauled off for repairs.

For decades, the technology did not exist to flag the problem before it led to such a system failure. Historically, operations analysts had to manually diagnose thousands of locomotive asset faults each day to assess whether components need repairs or replacing. This time-consuming process resulted in inefficient spending and unnecessary downtime.

With the influx of connected asset data and advanced analytics capabilities, that function can be dramatically optimized, with direct—and significant—impact on a railroad operation's bottom line. Uptake's platform addresses rail operators' most vexing questions: Does a locomotive need repairs before leaving the yard? When is the optimal time to pull locomotives off a track to send to the shop for service? How can teams get them back into operation faster?

A Platform Approach

In designing Uptake's technical architecture to address these most pressing pain points in rail and other major industries, we took advantage of advances in cloud computing, data science, and machine learning. We combined these technologies with the domain knowledge of major industry partners and collaborated with them to build and customize our platform for their industry.

This approach ensures that we are answering the right problems with the best data. For example, the design supports multi-tenancy, meaning a single instance of an application supports multiple customers, matching the model of our industry partners, where the main organization has multiple customers who have customers, and so on. By solving problems not just for one industry but for many at once, we apply patterns and solutions across verticals, creating approaches that are not necessarily intuitive, but which provide smarter solutions and immediate value.

To generate actionable insights and predictions from partner data, we have to accommodate, among other things, various data formats, types of storage, and different levels of access. During data acquisition, we gather data in batch or in real time as needed (via our REST API and ETL architecture) and store them in a secure cloud environment, with each partner's data in a separate data enclave. The mission-critical nature of the data that Uptake processes requires an approach to sensor-generated data security that is a leading model for this evolving space. This security model is based on leading practices and recognized industry standards, including the Cloud Security Alliance (CSA), NIST, and ISO, while applying these standards to new tools and approaches being used for large-scale predictive analytics.

We ingest diverse data at high speeds and then normalize the data for accuracy, efficiency, and quick access. The Uptake platform is capable of reducing multiple forms of data at global scale to its canonical form.

Since no single data store is best for the wide range of problems we solve, we built our platform on a polyglot data store capable of using more than one data type. For massive amounts of sensor data, for instance, we use a time series data store. To find relationships and map networks, we use a graph data store. To manage mission-critical data and transactions, we use a relational database. As we encounter more types of data, we will use the data store best suited to its structure. In this way, we ensure the fast and ready availability of data to our predictive engines and models.

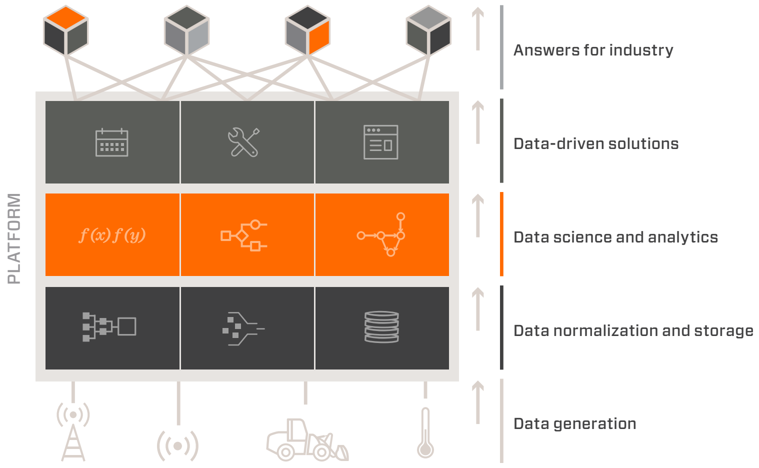

Our data science team has constructed modeling engines that drive intelligent recommendations and answers across industries. The scalability of the architecture supports large quantities of highly complex algorithms, while the data science engines are powerful enough to make real-time predictions based on live, streaming data. For added flexibility and efficiency, our data science models are integrated directly into our platform. While most companies implementing complex models in production have engineers recode data science models line-by-line, which wastes valuable time and resources, our model deployment merely requires answers to four questions: 1) When does the model run? 2) On what assets does this model apply? 3) What input data is needed? and 4) What is the output of the model?

We also designed our platform for continuous improvement in all aspects. For our data science models, we apply a "survival of the fittest" methodology to our algorithms so that the ones with the best results for a specific problem are used. The framework is language-agnostic, enabling the newest algorithms to be easily built without regard to implementation language. We also implement closed-loop feedback, feeding results of recommendations and actions back into the system. Thus, the platform gets smarter over time. Furthermore, we can apply learnings across industries, giving our partners insights that they ordinarily could not leverage.

By inserting the insights into human-driven operations and capturing the results, the models get smarter, and importantly, human "hands-on" knowledge becomes codified and preserved. This virtuous cycle of insight transforming effective action into better insight has the added benefit of empowering the training and performance of decision-makers and operators who may be new to the field.

Hands-on Insights Drive Results

Back to our example in rail, capitalizing on this "hands-on" institutional knowledge has been critical to building the best solution. Through our partnership approach, we have worked closely with leading manufacturers' experienced condition monitoring analysts who, until the development of the Uptake platform, manually assessed thousands of asset fault notifications each day. Their expertise and consistent feedback are directly reflected in solution features and continually updated as the software quickly learns and improves.

A key feature that has benefited from this critical feedback source is the Uptake platform's proprietary locomotive health score, which represents the probability that a critical fault will occur in the next two weeks. Uptake's model is at least three times more efficient than the industry standard procedure. This insight enables Uptake to deliver real-time assessments of locomotive health, drilling down to the condition of specific components. Railroads on the Uptake platform can now remotely troubleshoot problems, expedite maintenance by advanced ordering of parts, and clear shop time so that locomotives can quickly come in and out. This results in increased accuracy of repairs and improved cycle time at the shop, translating into more efficiency and productivity.

Further, because we calculate health scores of locomotives in real time based on the health of key components, analysts can better manage their fleets relative to the priority of the job and distance of the mission. This health score also identifies which subsystem of a locomotive needs attention, further reducing potentially costly shop time.

Uptake's platform has also contributed to improved fleet reliability through our Location Management tool, which geospatially tracks assets on a satellite image feed and provides profile and machine health records with a single click. Uptake's platform enables operators to make more informed maintenance decisions based on asset status, location, performance, and shop history. Ensuring that locomotives are properly maintained and correctly repaired reduces that industry's key reliability metric of FLY (Failures Per Locomotive Year). It is estimated that a FLY reduction of just 0.1 percent can translate into millions of dollars in savings for a typical U.S. railroad.

Bottom Line & Global Impact

As the data landscape continues to rapidly evolve, predictive analytics will become a necessary tool with far-reaching implications for an organization's optimization (and survival). We have already witnessed how data platforms have revolutionized consumer products such as fitness trackers that monitor exercise and accordingly recommend health regimens, hotel and restaurant reservation tools that forecast demand, or customized media playlists that predict genres for a user's tastes.

> "The sheer scale of these industries alone means the potential for new predictive analytics solutions can create hundreds of millions (or even billions) of dollars in new value and improved operations that, beyond the bottom line, stand to improve health and safety on a global level. **That is the promise of predictive analytics.**"

The largest, most pressing need for improvement lies within data-rich industries that help power the world. Water utilities, for example, can tap predictive analytics to better target infrastructural repairs to prevent water main breaks and floods, or to more effectively distribute chemicals to treat and purify their supply. Public safety agencies can better identify proactive opportunities to enhance patrol activities and serve as a source of help when and where it is most needed. The automotive industry can begin to improve road safety through monitoring driver behavior and noting risks for collisions and hazards, or predict just-in-time maintenance or part replacement. Health care providers can diagnose patients quicker, manage treatment and tests, and easily match specialists with the patients in most need.

The sheer scale of these industries alone means the po-

tential for new predictive analytics solutions can create hundreds of millions (or even billions) of dollars in new value and improved operations that, beyond the bottom line, stand to improve health and safety on a global level. That is the promise of predictive analytics. Q

***Brad Keywell** is the CEO and cofounder of Uptake Technologies, a predictive analytics SaaS platform provider that transforms data into actionable insight for productivity, efficiency, and operational safety across major industries. Among the previous ventures Keywell cofounded are Lightbank, Groupon (NASDAQ:GRPN), MediaOcean, and Echo Global Logistics (NASDAQ:ECHO). He is also founder and co-chairman of Chicago Ideas, an annual innovation platform for global thought leaders, and the former chairman of the Illinois Innovation Council. Keywell received a B.B.A. and his J.D. from the University of Michigan.*

References

1. Essick, Kristi. "The Internet of Everything - Four Technologies We'll Actually Use within Three Years." The Network: Cisco's Technology News Site (March 2014): https://newsroom.cisco.com/feature-content?type=webcontent&articleId=1373457.
2. IDC Digital Universe Study, The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things (April 2014).
3. Ovum Research, as cited in Geiger, Jakki. "How Much Does Bad Data Cost Your Business?" Informatica (April 2014): http://blogs.informatica.com/2014/04/16/do-you-know-how-much-bad-data-is-costing-your-business.

Unpowering the Internet of Things

by Joshua R. Smith

Will the Internet of Things (IoT) turn out to be the Internet of Dead Batteries (IoDB)? The IoT is envisioned as a vast collection of internet-connected, sensing-computing nodes distributed throughout our physical environment and must be powered somehow. If only batteries are used, then the IoDB will arrive as soon as they start failing en masse. To avoid this dismaying possibility, our team has been developing technologies to harvest power from both ambient and deliberately transmitted radio waves, and to communicate using many orders of magnitude less power than traditional radios. This will eventually allow us to create battery-free sensing nodes and further broaden the reach of IoT. In this article, we present a series of systems ranging from an accelerometer to a wireless microphone to a camera, all of which are powered by radio waves and contain no intrinsic sources of power.

The energy efficiency of microelectronics has improved by a factor of about one trillion (10^{12}) between 1940 and 2010[1,2]. This means that radio waves that we used to think of as carrying only information can now serve as sources of power capable of operating non-trivial electronic devices.

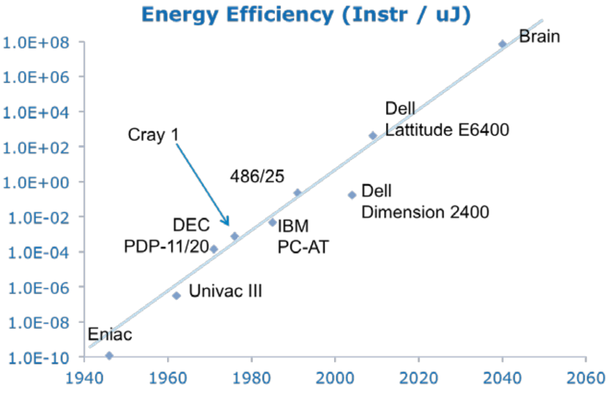

Figure 1 | The energy efficiency of microelectronics has improved by a trillion-fold since 1940[1].

Wireless Identification and Sensing Platform

The Wireless Identification and Sensing Platform (WISP) was the first UHF-powered, fully programmable microcontroller system[3,4]. WISPs are powered by 915 MHz UHF RFID readers and communicate with the reader by backscatter. The WISP platform was used to create the

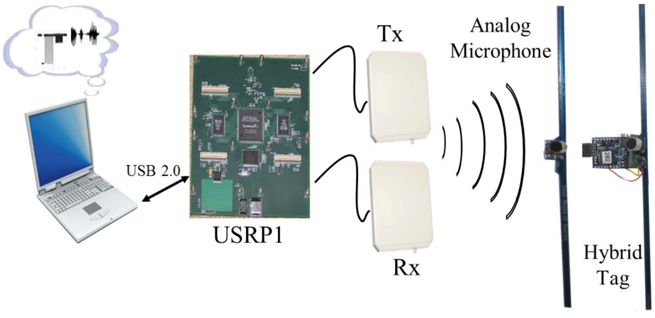

Figure 2 | Battery free microphone system. From left to right: Base station Computer; Universal Software Radio Peripheral (for RF and signal processing); Base station Transmit (Tx) and Receive (Rx) antennas; analog backscatter microphone tag; hybrid analog-digital backscatter microphone tag.

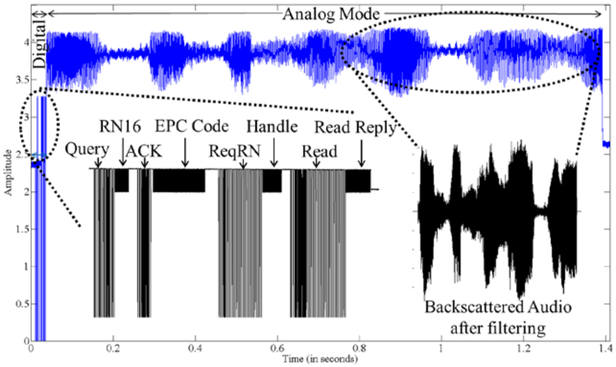

Figure 3 | Operation of hybrid analog-digital, battery-free microphone tag. The top trace is the signal observed at the base station. First, a two-way digital interaction between reader and microphone tag occurs, which allows the reader to identify and enable the correct sensor; then the tag enters analog mode and backscatters audio signals.

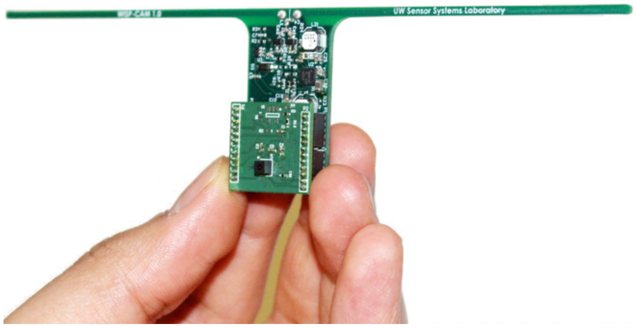

Figure 4 | WISPCam. The square daughterboard hosts the camera, which was originally produced for mobile phones. The main board behind the camera provides UHF power harvesting and bi-directional backscatter-based communication.

first UHF-powered accelerometer system[1] and the first UHF-powered strong cryptographic system[5].

More recently, we have been working on building RF-powered versions of more challenging and useful sensor systems, in particular microphones[6] and cameras[7].

Battery Free Microphone

The battery free microphone uses Analog Backscatter, in which the small charge output of the electret microphone actuates a transistor, which modulates the reflection coefficient of the analog backscatter tag's antenna. The reader collects the RF signals differentially reflected by the tag and reconstructs the audio stream. Figures 2-4 illustrate the analog backscatter microphone system.

RF-Powered and Read Camera

Recently we created the WISPCam, which we believe is the first UHF powered and read camera, meaning that the device is powered and read by a commercial off-the-shelf UHF RFID reader. The device benefits from small mobile phone-scale cameras, but mainly enabled by Ferroelectric RAM (FRAM), a new ultra-low-power non-volatile memory. The WISPCam charges a supercapacitor until it has sufficient energy (around 20mJ) to take a picture, stores it to FRAM, and then begins backscattering it. If the WISPCam runs out of energy before it finishes transmitting the data, it will sleep until it has accumulated enough energy to resume data transmission.

The original WISP consumed between 1-1000 μW. The power available for it to use is determined by distance from the RF source and as a result the WISP automatically adjusts its active duty cycle in order to match its power consumed to the power available. While sleeping, the WISP continues to harvest energy. By increasing the interval between wake events, the WISP can lower its average power consumption as much as desired. In practice this means that when a WISP is close to the RF source, it will typically perform hundreds of sense-compute-respond cycles per second; when it is far from the RF source, it may perform less than one per second.

Wireless Ambient Radio Power [WARP]

It turns out that approximately the same power levels can be collected from ambient RF sources such as TV towers[8], cell phone towers[9], or Wi-Fi access points[10].

Ambient Backscatter Communication

It is nice that ambient RF signals can be used to power IoT sensing devices, but what good is that if they are unable to communicate? With Ambient Backscatter Communication (ABC), we showed that the same ambient radio signals used for the power source can also serve as a communication "substrate" for IoT devices; they can communicate by selectively reflecting the same pre-existing ambient radio waves that are their power source[11] This allows nearby devices to communicate with one another. By introducing coding gain, the communication range can be increased by two orders of magnitude[12], which can enable long-range reading of such sensors.

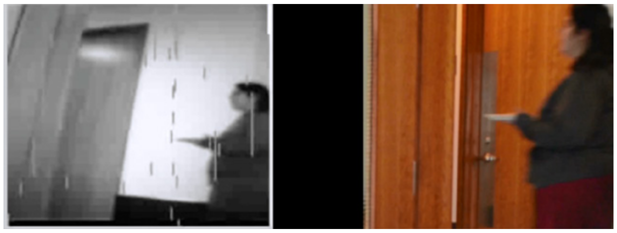

Figure 5 | Left: An image captured by a motion-triggered WISPCam. Right: ground truth image captured by a conventional camera. In this WISPCam variant, a passive IR sensor is used to trigger image capture when the door is opened.

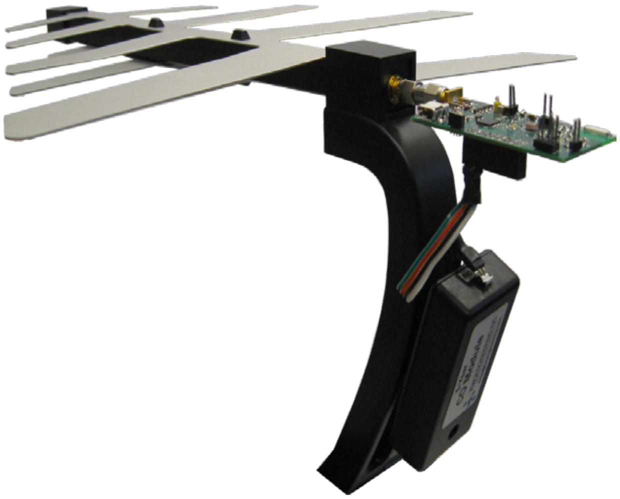

Figure 6 | Ambient RF Harvester, coupled to an electrolytic gas sensor[13].

Passive Wi-Fi

With the Passive Wi-Fi project, we showed that it is possible for an ultra-low power device to use backscatter (reflection of pre-existing radio waves) to generate ordinary, standards-compliant 802.11b Wi-Fi packets that can be received by any ordinary Wi-Fi receiver with no special backscatter receiver required. With this system, ultra-low power IoT endpoints can communicate with the ordinary, pre-existing internet with no need for an IoT gateway.

The Future

Using directed and ambient RF signals as a power source, we have demonstrated a sequence of increasingly capable battery-free sensors ranging from a simple accelerometer to a wireless camera. There are challenges to overcome, such as designing more effective coding schemes for ambient backscatter communication, and implementing long-range ambient backscatter reader basestations, but our experience so far suggests that battery-free, RF-powered sensors are feasible now, and becoming increasingly practical.

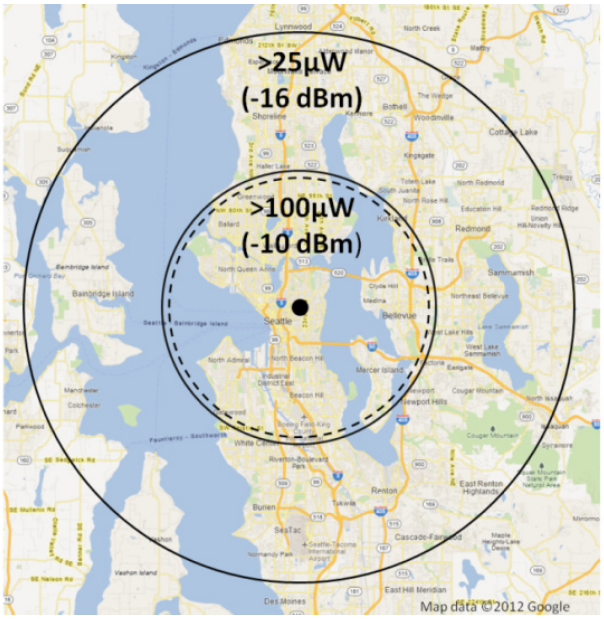

Figure 7 | Map of Seattle, showing power available from a particular TV tower.

***Joshua R. Smith** is an Associate Professor in the departments of Computer Science & Engineering and Electrical Engineering at the University of Washington, where he leads Sensor Systems Laboratory. He was named an Allen Distinguished Investigator by the Paul G. Allen Family Foundation and he is a Thrust Leader in the NSF Engineering Research Center on Sensorimotor Neural Engineering (CSNE). His research focuses on inventing new sensor systems, devising new ways to power them, and developing algorithms for using them. In addition to intelligence, this research has applications in the domains of implanted medical devices, robotics, and ubiquitous computing. Smith has co-founded two companies to commercialize his research: Jeeva Wireless Inc. and Wibotic Inc.*

References

1. Smith, J.R. Range scaling of wirelessly powered sensor systems. Wirelessly powered sensor networks and computational RFID (JR Smith Ed.). pp. 3-12. Springer SBM 2013.
2. Koomey, J.G. and S. Berard, M. Sanchez, H. Wong. Implications of historical trends in the electrical efficiency of computing. Annals of the History of Computing. IEEE, 33(3):46–54. March 2011.
3. Smith, JR and AP Sample, PS Powledge, S. Roy, A. Mamishev. A wirelessly-powered platform for sensing and computation. UbiComp 2006: Ubiquitous Computing. pp. 495-506. 2006.
4. Sample, AP and DJ Yeager, PS Powledge, AV Mamishev, JR Smith. Design of an RFID-based battery-free programmable sensing platform. Instrumentation and Measurement. IEEE Transactions on 57 (11). pp. 2608-2615. 2008.
5. Chae, HJ and M. Salajegheh, DJ Yeager, JR Smith, K. Fu. Maximalist cryptography and computation on the WISP UHF RFID tag. Wirelessly Powered Sensor Networks and Computational RFID (Smith Ed.). pp. 175-187. 2013.
6. Talla, V. and M. Buettner, D. Wetherall, JR Smith. Hybrid Analog-Digital Backscatter Platform for High Data Rate, Battery-Free Sensing. IEEE Topical Meeting on Wireless Sensors and Sensor Networks (WiSNET). 2013.
7. Naderiparizi, S. and AN Parks, Z. Kapetanovic, B. Ransford, JR Smith. WISPCam: A Battery-Free RFID Camera. IEEE RFID 2015.
8. Sample, A. and JR Smith. Experimental results with two wireless power transfer systems. Radio and Wireless Symposium, 2009. RWS'09. IEEE. pp. 16-18. 2009.
9. Parks, AN and AP Sample, Y. Zhao, JR Smith. A Wireless Sensing Platform Utilizing Ambient RF Energy. IEEE Topical Meeting on Wireless Sensors and Sensor Networks (WiSNET). 2013.
10. Talla, V. and B Kellogg, B. Ransford, S. Naderiparizi, S. Gollakota, JR Smith. Powering the Next Billion Devices with Wi-Fi. ACM CoNext 2015.
11. Liu, V. and A Parks, V. Talla, S. Gollakota, D. Wetherall, JR Smith. Ambient backscatter: wireless communication out of thin air. Proceedings of the ACM SIGCOMM 2013 conference, pp. 39-50. 2013.
12. Parks, A.N. and A. Liu, S. Gollakota, J.R. Smith. Turbocharging Ambient Backscatter Communication. Proceedings of the ACM SIGCOMM conference, 2014.
13. Carter, M.T. and J.R. Stetter, J.R. Smith, A.N. Parks, Y. Zhao, M.W. Findlay, V. Patel. Printed Low Power Amperometric Gas Sensor Employing RF Energy Harvesting. The Electrochemical Society. January 2012.

Acknowledgments

This survey paper describes joint work with collaborators Shyam Gollakota, as well as students and former students Alanson Sample, Vamsi Talla, Aaron Parks, Saman Naderiparizi, and Zerina Kapetanovic.

This work is funded by Google Faculty Research Awards, a Google Ph.D. Fellowship, a Qualcomm Innovation Fellowship, the Intel Science and Technology Center for Pervasive Computing, and NSF Grants CNS-1305072, EEC-1028725, CNS-1407583.

From the IQT Portfolio

The *IQT Quarterly* examines trends and advances in technology. IQT has made a number of investments in IoT technologies and several companies in the IQT portfolio are garnering attention for their unique solutions.

Mocana

Mocana provides a device-independent security platform that secures all aspects of mobile and smart connected devices, as well as the apps and services that run on them. Mocana recently announced the integration of its Security of Things Platform with Schneider Electric's energy management devices and Latronix's IoT networking solutions. Mocana is located in San Francisco and became an IQT portfolio company in March 2012. **www.mocana.com**

Orion Labs

Orion Labs offers a small wearable device that enables push-to-talk communication over any available network. The company recently launched its first product, Onyx, which pairs with a mobile app and uses a smartphone's data or WiFi connection. Orion Labs is based in San Francisco and joined the IQT portfolio in September 2009. **www.orionlabs.io**

PsiKick

PsiKick is redefining ultra-low power wireless sensing devices – developing the lowest power devices of their kind in the world. PsiKick's technology provides the basis for this vision of ubiquitous computing and, eventually, a true Internet of Things (IoT) by embedding self-powered awareness into any device, object, building, structure, or environment. The company hopes to deliver a full wireless system of batteryless sensor nodes by the end of 2016. PsiKick is based in Charlottesville, Va. and joined the IQT portfolio in June 2015. **www.psikick.com**

PubNub

The PubNub data stream network provides the cloud infrastructure and key building blocks for real-time apps that scale globally to any device. PubNub enables real-time experiences like live dashboards and streams, presence, collaboration, second-screen synchronization, machine-to-machine signaling, and more. The company was recently featured in *Forbes* and the *San Francisco Business Times* for a top 'office hack' among tech startups: PubNub engineers integrated a coffee maker with a Raspberry Pi device to notify a user when coffee needs replenishment. PubNub is located in San Francisco and became an IQT portfolio company in March 2014. **www.pubnub.com**

IQT®

IN·Q·TEL