



Enabling the Internet of Things

Roy Want, Bill N. Schilit, and Scott Jenson, Google

Merging the virtual World Wide Web with nearby physical devices that are part of the Internet of Things gives anyone with a mobile device and the appropriate authorization the power to monitor or control anything.

The Internet of Things (IoT) paradigm enables interconnectedness among devices—anytime, anywhere on the planet—providing the Internet’s advantages in all aspects of daily life. Analysts predict that the IoT will comprise up to 26 billion interconnected devices by 2020, a 30-fold increase from 2009 (www.gartner.com/newsroom/id/2636073).

The conventional Internet has proved valuable in almost all endeavors by giving people the ability to interact with global information and services. The majority of this interaction happens through the World Wide Web, with client computers running a browser and communicating with cloud-based servers. However, the Internet is not limited to the Web: a wide diversity of other protocols are employed to make use of global Internet connectivity. The IoT is considered to be the next logical evolution, providing extensive services in manufacturing, smart grids, security, healthcare, automotive engineering, education, and consumer electronics. Many of these systems already have a Web presence but use protocols that are largely Web independent.

Practical issues with the IoT vision must be addressed,

including how to handle dramatic increases in network scale and how to determine device proximity, sometimes referred to as localized scalability.¹ In an IoT world, preferentially discovering things nearby and letting users interact with them is a powerful mechanism for overcoming a global network’s scale and complexity. Other important IoT enablers are peer-to-peer connections, low-latency real-time interaction, and integration of devices that have little or no processing capability.

THE IOT VISION

The Web provides an important interaction model for the IoT by letting users get device-related information and in some cases control their devices through the ubiquitous Web browser. The conventional Web is a convenience we enjoy as we search for information, respond to email, shop, and engage in social networking; the IoT would expand these capabilities to include interactions with a wide spectrum of appliances and electronic devices that are already ubiquitous in the early 21st century.² We refer to devices that are part of the IoT and directly accessed, monitored, or controlled by Web technologies

as the Physical Web: Physical Web = Web technology + IoT.

Identifiers are the key to enabling any kind of interaction among devices. From an IoT perspective, IPv6's 128-bit addresses serve as identifiers for a global network of devices. Alternatively, Uniform Resource Identifiers (URIs), which include both locators and names, provide a higher-level concept that bridges those devices to existing Web technology. The Uniform Resource Locator (URL) is used in conjunction with a Distributed Name Service (DNS) to route and connect to services. Uniform Resource Names (URNs), such as globally unique IDs, are resolved by scheme-specific methods. A distinguishing aspect of the Physical Web is to consider URIs as the primary identifier.³

Many researchers and practitioners in this field, including the authors of this article, expand the IoT definition to include enabling an Internet presence for any person, place, or thing on the planet, thereby pushing our notion of the Physical Web beyond smart devices. Clearly, an Internet presence cannot occur without processing and networking, so instead of providing them directly, an Internet service can provide information and perform actions via other nearby devices serving as a gateway to that proxy service.^{4,5}

Gateway devices will enable billions of people, places, and things to participate in the IoT—most people today already carry one. The smartphone, the most popular computing device of all time, with more than 1 billion users (www.idc.com/prod_serv/smartphone-os-market-share.jsp), is well equipped to serve as this pervasive portal.

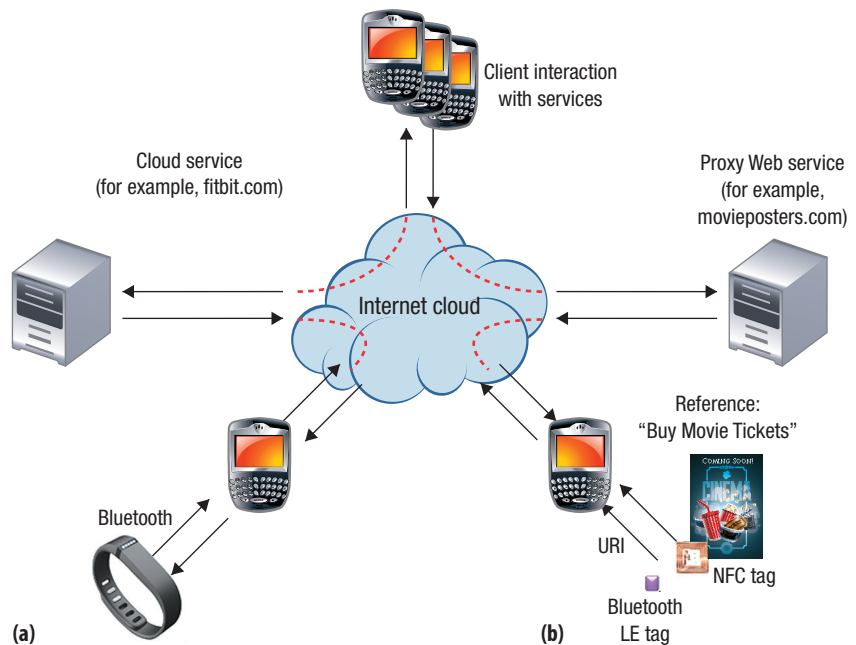


FIGURE 1. Two methods for a smartphone to interact with the Internet of Things: (a) direct and (b) proxy. Networked computers can participate in the IoT—passive objects can, too. LE, low energy; NFC, near-field communication; URI, Uniform Resource Identifier.

Figure 1 shows the two distinct interaction modes that smartphones can enable in the IoT. Through *direct interaction*, a smartphone can query the state of an IoT device in its proximity and then provide a bridge between low-level peer-to-peer protocols, such as Bluetooth or Wi-Fi, and Internet protocols, such as HTTP and TCP. One example is the Fitbit fitness monitor, which uploads a user's step count through his or her phone over a 4G network to the user's account in the cloud. Through *proxy interaction*, mobile users who happen to be near an IoT-enabled object or device can look up associated information published by interested parties through a Web service using their smartphone, just as they would when performing a Web search. One example is a movie poster that enables nearby people to automatically access a webpage on their smartphone and buy electronic tickets online.

DOES THE IOT ALREADY EXIST?

The IoT is a popular buzzword in the computing industry; it appears in the marketing campaigns of

major networking companies such as Cisco and microprocessor giants such as Intel. It even serves as the title or theme of conferences, such as the "Internet of Things" World Forum (<http://iotinternetofthingsconference.com>).

However, the phrase represents ideas that have existed since the beginning of the Web or been written about in whitepapers from well-known research laboratories such as (Xerox) PARC and HP Labs. So why isn't the IoT a standard part of the way we do business today? Why is it still the subject of speculation and vision statements in keynote addresses at well-known computer industry events such as the annual Consumer Electronics Show?

The answer appears to be that the IoT exists for a small number of technologies that have the ingredients for a successful business case. In general, these early systems have tended to be closed ecosystems, using private APIs and locking up the data. This is counter to the spirit of open systems at the heart of the original Internet standards, reflecting instead the more recent commercial successes of




FIGURE 2. Various forms of electronic tags support the Physical Web (all about the size of a quarter): (a) a near-field communication (NFC) tag; (b) a quick response (QR) code; and (c) a Bluetooth low energy (BLE) tag. However, it is not clear which technology—each with its own affordances and problems—will become the primary IoT enabler.

proprietary business entities such as Apple’s App Store and Facebook. You can actually buy home automation systems that connect to the Internet through your home’s Wi-Fi. These systems are usually built with a bridge that controls the automation components through proprietary protocols on one side and communicates with open protocols to a proprietary Web service on the other. Users can then employ desktop computers or smartphones as a client to control their home by interacting with the Internet service, effectively providing user interface hardware at no cost to the IoT device manufacturer.

A significant hurdle to fully realizing the IoT relates to scale—specifically, expanding the Internet to IoT scale means that the address space for the Internet will need to increase by several orders of magnitude. Therefore, another requirement for supporting the IoT is a larger device address space than that provided by IPv4. To enable this kind of expansion, the Internet Engineering Task Force (IETF) has been working on the IPv6 standard for some time. When the transition is complete, the address space will be large enough to support every object on the planet, enabling embedded computers of all sizes to be easily integrated into the Internet. However, a large percentage of the objects in the IoT will not be suitable for direct

wired or wireless connection to the Internet, falling into the class of passive devices. For these objects, a tag, smartphone, and proxy Web service is needed to provide users with the object’s Web presence. Of all the visionary ideas around the IoT, this one has made the least progress to date.

ENABLING TECHNOLOGIES FOR THE IOT

As Figure 2 indicates, tagging an object to reference a proxy Web service can be achieved through a variety of technologies, but the early primary contenders have had issues that hindered their adoption.

RFID and near-field communication

In the early 2000s, RFID was considered one of the most likely technologies to accelerate the formation of the IoT.⁶ A new UHF RFID tag standard was developed by EPC Global (<http://epcglobal.org>), with a goal of further automating retail transactions and replacing barcodes with a tag that was machine readable at a distance of up to 10 feet. But after several trials with leading vendors such as Walmart and Tesco, the EPC standard met with limited uptake principally because, in practice, a significant number of tags were undetectable due to factors in retail environments such as poor product/tag orientation and the presence of materials

that interfered with the wireless identification process.

The latest opportunity for RFID technology is in the form of near-field communication (NFC) as a support for electronic payments. Although only a small number of smartphone products include NFC transceivers, the potential for this capability to propagate to all future smartphones is high. It would also enable phones to read passive NFC tags that can store a URI, while still being cheap, small, thin, and attachable to almost anything. In September 2014, Apple announced the iPhone 6 would include NFC support for ApplePay. With Apple’s significant smartphone market share in the US, this move could influence other handset manufacturers to follow suit, pushing NFC into becoming a key IoT enabler.

Optical tags and quick response codes

Another contender for low-cost tagging is the optical or printed tag—in particular, the most popular 2D optical standard, the quick response (QR) code.⁷ The success of the QR code standard is directly related to the ubiquity of its reader, an application of the high-resolution camera found in all modern smartphones. A QR code is extracted and decoded from a scene using image-processing techniques yielding a number, text, or URI. QR codes are already printed on many products,

IN THE PHYSICAL WEB, PEOPLE, PLACES, AND THINGS HAVE WEBPAGES TO PROVIDE INFORMATION AND MECHANISMS FOR USER INTERACTION.

including newspapers, magazines, billboards, and coupons; they even appear on prime-time television ads.

However, in practice, many QR advertising campaigns result in a poor customer response. The reasons relate to the requirement that a preinstalled application is required to read QR codes—which can be a barrier for some users—as well as difficulty in positioning the phone so that the camera can focus and accurately decode the image. Some advertisers also feel that a visible QR code spoils the aesthetics of their campaigns (<https://www.techdirt.com/blog/wireless/articles/20120307/06130018010/qr-codes-ugly-overused-doomed.shtml>).

Bluetooth low energy

One of the more promising new technologies in the device tagging space is Bluetooth low energy (BLE),⁸ part of the Bluetooth v4.0 standard (Bluetooth Smart) and adopted by the Bluetooth Special Interest Group in 2010. Consequently, Bluetooth silicon vendors have included BLE in their latest chipsets, and all smartphones released in the last few years have BLE hardware, with various levels of capability depending on operating system support.

Bluetooth silicon can be pared down to only include the BLE aspects of the standard, removing the need for compatibility with classic Bluetooth. This results in a small, low-cost silicon implementation that can be used as a low-power electronic tag. Tags based on BLE can signal their presence by transmitting an advertising packet once per second at a power budget that enables them to operate for up to one year on a lithium coin cell battery (about the size of a US quarter, with 240-mAh capacity).

This new technology standard, along with the availability of inexpensive BLE tag hardware and tag readers already integrated with smartphone hardware, has been a considerable catalyst in this space. Many established computer companies, and a significant number of startups, are experimenting with products and business opportunities associated with these tags. As with the other tagging technology, this is also an enabler for the IoT, but with more opportunities for ubiquitous deployment, higher-accuracy tag reads, and the ability to blend in invisibly with a product.

THE PHYSICAL WEB

In the Physical Web, people, places, and things have webpages to provide information and mechanisms for user interaction. The notion of open Web technologies as the bridge to the physical is not new: access points, routers, solar panels, electricity meters, and coffee shops have Web landing pages, for example. However, it is the breadth and depth of the stack surrounding the Web that make this an appealing vision for the IoT's evolution. To be sure, HTTP will not be an exclusive protocol for communication with things in the same way that it is not an exclusive protocol for the Internet—there are plenty of use cases where Web protocols do not have the desired properties, such as the Real-Time Streaming Protocol (RTSP).

Like the conventional Web search engine to which you submit a query and it returns text snippets and links to relevant “things,” the Physical Web would return search results. However, because the IoT is the world that we can see, hear, and touch, search results would not only be ordered by conventional ranking algorithms but also by

proximity, and thus results could be shown as lists, enhanced maps, or floor plans. Searching the physical Web at home would bring up thermostats, DVRs, TVs, home audio systems, routers, and water and electrical meters as results. You might see a snippet and a link to the manual for a nearby microwave oven, along with other links that control or provide information about devices in your house. There would be a lot to show, but no more than the results of any Web search today; search engines are good at ranking and providing the most relevant items first.

Webpages are a great technology for human-to-machine (H2M) interaction, but many use cases for the IoT are machine to machine (M2M). One solution that has already had great success in combining human and machine-readable content in open Web technologies is the use of structured data embedded by webmasters into their pages. The data formats used by Schema.org and others let user agents and cloud services act intelligently, parsing data for events, organizations, people, places, products, reviews, and so on and acting on them either interactively or proactively. Structured data could also lead to more uniform user interfaces across devices, so that when users learn one interface, they do not have to relearn it for another device—for example, setting the time on an appliance that has a clock.

Open Web technologies, including HTML, Ajax, HTTPS, and OpenID, and structured data apply equally well to the IoT. However the open Web does not have an effective mechanism for locating objects in the physical world. One approach that looks particularly appealing is the use of radio beacons that broadcast URLs at very low power and over a small geographic radius.

OUTLOOK




FIGURE 3. Alternative approaches for the IoT. (a) Peer-to-peer, with gateway to an IoT service); and (b) centralized IoT service. Because the IoT will connect many devices with constrained networking capabilities, P2P networking may become more prevalent and as common as cloud-based solutions are now.

One of the first projects to promote this idea was HP's Cooltown,³ which used infrared beacons to transmit URLs. More recently, BLE provides a similar low-power beacon mechanism that can also integrate a URL emitted in short periodic advertisement packets (www.uribeacon.org).

One way to bridge the gap between the physical and virtual is to attach beacons to all our objects. These beacons would broadcast a URL along with other information to help with ranging. It sounds like a huge investment, but beacons currently cost less than US\$5, and the price is likely to drop, thus future manufactured "smart" objects are likely to integrate this capability. By utilizing proximity URL beacons, the potential for interacting with the Physical Web is not only more practical, but has greater utility than both NFC and QR codes.

CLOUD COMMUNICATION VERSUS PEER TO PEER

One of the long debates in computer science has been whether to build centralized systems or to make them fully distributed. Much of the Web

today takes a centralized approach for its services, but it is not clear that this makes sense for much of the IoT. Figure 3 compares the approaches, with cloud-based services representing the centralized paradigm.

Benefits of cloud computing

The computing world has shifted paradigms several times, from the centralized mainframe computer to the decentralized PC running standalone applications back to today's centralized cloud services. The computing industry is gravitating toward more centralized cloud services primarily because they are easier to manage. Advantages include the economics of scale when building datacenters, automatic backup of all data, and enforced physical security. However, modern client devices are both capable and flexible. Laptops utilize high-performance multicore technology, and even smartphones contain powerful processors. We have the option of running simple clients connected to powerful cloud services or powerful local apps that run on their own. The decision comes down to our tolerance for tradeoffs in latency,

security, privacy, and cost. If interaction latency and connectivity are not a problem, cloud computing is an attractive paradigm. In recent times, the detrimental effect of malware on home computers has made the cloud even more attractive.

Based on these observations, it seems like a good idea for the IoT architecture to register every device with a cloud service and communicate with that service alone. Users or other computers would then interact with this service to determine the device's status or control its behavior. However, there are other factors to consider.

Benefits of peer-to-peer

Although the cloud model is clean and straightforward, the requirements for full Internet communication might be too costly or burdensome for what in many cases are simple low-performance devices. In practice, any hardware that can connect directly to the Internet would require a physical Ethernet, Wi-Fi radio, or cellular modem, all of which elevate the device's cost and power consumption. In practice, it might be better to have one bridging device that supports Wi-Fi and enables simple peripheral IoT devices to talk to the bridge. Many low-performance communication standards such as ZigBee, 6LoWPAN, Bluetooth Classic, and BLE have evolved to fill that gap—with no clear winner. As with the opportunity for tagging devices, the new BLE standard could well dominate this low-end space, but it will take time for this to play out.

IoT devices, unlike the traditional Internet, benefit from the concept of proximity as we described when introducing the Physical Web. However, it is not just the knowledge of nearby objects that is useful: colocated devices have the opportunity to cooperate

THE LONG TAIL OF THE MANY POTENTIAL USES OF THE IOT CHALLENGES THE SCOPE AND BREADTH OF TODAY'S SMARTPHONE APPS.

with each other in real time and fulfill a task that would not be possible with any single device.⁹ A simple example is the sharing of peripherals. A computer that finds itself close to another device with a bigger and better screen could wirelessly share it, for example, to take advantage of its display when playing a movie. Likewise, a computer could utilize a component that it does not have if it is wirelessly discovered on another nearby device—examples include a sensor, camera, or mouse.

Proximate sharing will be an important aspect of the IoT, but this will not be easy to accomplish. Devices need to discover each other, trust each other, and then make a connection. Furthermore, sharing is achieved by protocols and data formats that need to be standardized and supported by both ends of the connection. In a world in which there are many more commercial players than in the early days of the Internet, this becomes a difficult proposition. The traditional Internet has evolved a core set of standardized protocols, but in the undefined IoT world, many standards and proprietary solutions are still up in the air. To illustrate how difficult this can be, consider the Digital Living Network Alliance (DLNA; <http://dlna.org>), which was established to enable multivendor consumer electronics to discover each other and share content and services. To date, DLNA has only been integrated with a small percentage of networked products because, in practice, larger companies are driven by the financial rewards of dominating a market with their own proprietary ecosystem.

Hybrid IoT solutions

Innovative solutions to latency problems associated with existing cloud-based networking attempt to combine

various ideas. One approach, called edge computing, moves some of the cloud processing closer to devices that require real-time interaction, thus reducing the number of network hops and hence latency. An example of edge computing is the cloudlet paradigm,¹⁰ which provides a means of rapidly enabling real-time services in the fixed infrastructure to be used by mobile devices, such as smartphones and wearables, just one network hop away using Wi-Fi. This is achieved through a virtual machine (VM) running on a powerful nearby workstation, and dynamically provisioned with software and services customized for that application. When the task is complete, the resources can then be freed up, allowing a new VM, or multiple VMs, to be instantiated on the cloudlet. A cloudlet lets IoT devices interact in real time with cloud-like services, even though they are far removed from a datacenter.

WHAT THE IOT MEANS FOR APPS

Today, any new IoT product almost always comes with a smartphone app to control it. This is a consequence of two dominant forces: native smartphone apps offer the only practical means to access and communicate with smart devices, and there has been so much recent financial success with mobile apps that they are now expected. People assume that they will always use apps for every possible interactive experience with IoT devices.

This approach unfortunately creates problems as the IoT landscape grows and matures. Although it is easy to imagine a phone with a few dozen applications on it, things become more problematic with millions of IoT devices. In the very near future, you will

likely pass thousands of smart devices every day, each one capable of interaction. It does not scale to have to install an app before using each one. You will also likely want to delete apps as you install others, because there will be many devices that you will interact with only once.

This becomes even more evident when you factor in how smart devices will have a significant long-tail effect.¹¹ Instead of a few big apps that you use every day for many tasks, you will have a huge range of small-device apps that offer the tiniest of interactions, such as just controlling an on/off switch. In fact, many devices will eschew interactivity all together and only offer a snippet of data, for example, when the next bus is arriving. The long tail of the many potential uses of the IoT challenges the scope and breadth of today's smartphone apps, requiring only micro-interactivity or micro-information.

To enable this new Physical Web for the IoT, we need a way for any user, with any smartphone or tablet, to walk up to any IoT device and interact with it (without a specialized app). We need a richer extension of today's Web, allowing each smart device to wirelessly broadcast a URL to its surroundings. Proximity is the context that can be used to filter this to a tractable number, allowing any smart device to list and then interact with other nearby smart devices. This is the basis of a discovery service that is the best of today's native and Web apps, creating a new platform, a type of interactive lingua franca allowing devices of any type to offer data and interactivity to any other device.

The current model of native apps, while quite popular, is not up to the task of supporting the many and

varied use cases that the IoT will bring about. We need to extend well-known systems, such as the Web, to allow a significantly easier and more lightweight approach for enabling devices to interact with one another.

Context sensing

One of the ways an application can perform more effectively is through context awareness. Sensing what is around a host device (and its user) and the context in which it is used allows

a display, actuator, or switch can be controlled from a browser or through a Web service, thus making it easy to integrate related information into control decisions. For example, emerging Web-connected irrigation systems might provide an interface for specifying the plants in your garden and use Web services to determine expert watering recommendations.

THERE IS A REASON PEOPLE ARE EXCITED ABOUT THE IOT: IT FEELS LIKE A BIG OPPORTUNITY TO IMPROVE HOW WE DESIGN AND BUILD PRODUCTS.

Privacy and security

One of the primary challenges for the future will be avoiding the darker consequences of a world with globally connected devices. The Physical Web could enable hackers to control our devices unless precautions are taken. The conventional Web already has security measures in place that can be applied to the Physical Web, but it is unclear if these will be suitable or even adequate for all IoT applications. In addition to hacking, social threats can result when knowledge is leaked in unexpected ways. For example, knowledge that a house is in an energy-saving mode could be a good indication that nobody is home and thus invite a burglar. These challenges will become more pressing as use of the Physical Web continues to grow.

FUTURE OPPORTUNITIES AND CHALLENGES

While the IoT applications will be varied and difficult to predict, some clear opportunities will arise for ubiquitous information gathering, context sensing, and control, which will likely be key enabling building blocks. However, some of the real challenges will be in the areas of privacy and security.

Ubiquitous information

Users often need help when they encounter an unfamiliar device. Traditional products were sold with a user manual—later, they might have included an informational CD. More recently, a customer support URL is provided that points to online documents, and even in the event that this is lost, a simple Web search for the model number will usually locate the online manual. In the IoT future, this problem can be solved by the Physical Web enabling all products to transmit a wireless machine-readable URL that can be received by a nearby smartphone or tablet with little effort from the user.

an app to adapt how data is presented and filtered; consider, for example, a local map application that automatically shows the user's current position. The more context a device has, the more likely it can automatically provide a user with the required information. However, the worst thing a context-aware application can do is infer something on behalf of a user that turns out to be wrong. The future IoT will allow information to be collected from nearby physical sensors and Web services and then shared between devices on a scale that has not been possible before. As a result, the Physical Web will lower uncertainty and improve accuracy. A consequence is that future "smart devices" will become much smarter.

Actions and control

As a complement to sensing, the IoT offers us a way to control the physical world through displays, actuators, and switches. Many modern systems benefit from remote control because it simplifies physical interaction design and extends capabilities. In the Physical Web paradigm, anything with


Merging the virtual Web with the IoT will enable *really* smart devices, providing smarter automation and services. Associating a Web URI with every person, place, or thing forms the basic mechanism for bridging these two technologies. A key enabler is the ability to discover proximate objects, and BLE beacons containing URIs are a useful tool. Web browsers can readily display Web information along with descriptions of nearby IoT objects; they are a promising familiar on-ramp to the IoT.

ABOUT THE AUTHORS

ROY WANT is a research scientist at Google. His research interests include mobile and ubiquitous computing. Want received a PhD in computer science from Cambridge University. He is an ACM and IEEE Fellow. Contact him at roywant@gmail.com or via www.roywant.com/cs.

BILL N. SCHILIT is a research scientist at Google. His research focuses on smart personal and mobile technologies supporting knowledge work. Schilit received a PhD in computer science from Columbia University. He is an IEEE Fellow, a member of the IEEE Computer Society and ACM, and associate EIC of *Computer*. Contact him at schilit@gmail.com.

SCOTT JENSON is a user experience strategist at Google. He has been in the design business for over 30 years, working on projects such as Mac System 7, Apple's User Interface Guidelines, the Newton, Google Mobile Maps, and now the Physical Web. Contact him at scott@jenson.org.

There is a reason people are excited about the IoT: it feels like a big opportunity to improve how we design and build products. Making the IoT as accessible and useful as the Web is likely an even bigger opportunity. 

REFERENCES

1. M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *IEEE Personal Comm.*, vol. 8, no. 4, 2001, pp. 10–17.
2. M. Weiser, "The Computer for the 21st Century," *Scientific Am.*, vol. 265, no. 3, Sept 1991; www.scientificamerican.com/article/the-computer-for-the-21st-century/.
3. T. Kindberg et al., "People Places and Things: Web Presence for the Real World," *ACM J. Mobile Networks and Applications*, vol. 7, no. 5, 2002, pp. 365–376.
4. R. Want et al., "Bridging the Physical and Virtual Worlds with Electronic Tags," *Proc. ACM SIGCHI*, 1999, pp. 370–377.
5. S. Jain et al., "Exploiting Mobility for Energy Efficient Data Collection in Wireless Sensor Networks," *Mobile Networks and Applications*, vol. 11, no. 3, 2006, pp. 327–339.
6. R. Want, "RFID: The Key to Automating Everything," *Scientific Am.*, Jan. 2004, pp. 56–65.
7. H. Kato and K.T. Tan, "Pervasive 2D Barcodes for Camera Phone Applications," *IEEE Pervasive Computing*, vol. 6, no. 4, 2007, pp. 76–85.
8. R. Heydon, *Bluetooth Low Energy*, Prentice Hall, 2013.
9. B.N. Schilit and U. Sengupta, "Device Ensembles," *Computer*, vol. 37, no. 12, 2004, pp. 56–64.
10. M. Satyanarayanan et al., "The Case for VM-Based Cloudlets in Mobile Computing," *IEEE Pervasive Computing*, vol. 8, no. 4, 2009, pp. 14–23.
11. C. Anderson, "The Long Tail," *Wired*, vol. 12, no. 10, 2004; <http://archive.wired.com/wired/archive/12.10/tail.html>.




Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

Software

On Computing

podcast

www.computer.org/oncomputing



with
**GRADY
BOOCH**



IEEE  computer society