Spring 2016 Volume 13, Number 1

iHOMES and BUILDINGS

THE MAGAZINE OF THE CONTINENTAL AUTOMATED BUILDINGS ASSOCIATION

Climate Change Driving Intelligent Building Adoption

New Year's Resolutions for Building Owners

Strong Partnerships Will Accelerate Mass Market IoT Adoption

What is the Internet of Things?

Project-Haystack: A CABA White Paper

How safe are home security systems?

CABA™ Connect to what's next™

www.caba.org/ihomesandbuildings

iHOMES and BUILDINGS

THE MAGAZINE OF THE CONTINENTAL AUTOMATED BUILDINGS ASSOCIATION

Spring 2016, Volume 13, Number 1

Contents

Ken Wacks' Perspectives

What is the Internet of Things?

By Ken Wacks

I first heard the phrase "Internet of Things" more than 15 years ago while consulting at the Massachusetts Institute of Technology (MIT). I was helping to develop the MIT Home of the Future, a project at the MIT Media Lab. We were searching for a networking technology to deploy sensors in a house designed to monitor daily activities of a volunteer as we investigated assisted living at home. I wrote about this project for the *CABA Home and Building Automation Quarterly*, the predecessor to *iHomes and Buildings*, in the winter of 1999. We examined research into techniques for embedding communication interfaces into sensors, so sensor data could be collected efficiently. The hope was that sensors for temperature, light level, movement, and mechanical motion (such as doors opened or faucets turned) could be deployed and networked inexpensively.

Professor Neil Gershenfeld of the MIT Media Lab introduced the term "Internet of Things" in a *Scientific American* article from 2004.[1] He predicted that the Internet of Things concept could enable low-cost sensors to be embedded into buildings to monitor the integrity of the infrastructure, such as the steel beams.

Meanwhile, standards for home and building systems were being specified and products were being introduced. These systems consist of interconnected sensors, actuators, controllers, and user interfaces to support applications such as lighting and energy management, as illustrated in Figure 1. These system components constitute a network of communicating things, but do not usually use the Internet for communications. Thus, CABA members are very familiar with a network of things, even if we have not used the phrase "Internet of Things."

The things

The fundamental tenet of automation is communications among devices without a human in the loop. In industrial control systems this is called "telemetry" from the Greek word meaning measurement at a distance. Typically a sensor measures a physical parameter, encodes the measured value into an electrical signal, and sends the signal on a wire or via radio to a distant receiver. At the receiver the sensor data are decoded and may be used by a controller to manage a system. For example, the sensor might be a thermostat that measures temperature and sends an electrical representation to equipment for heating, cooling, or ventilation.

The engineering challenges for a sensor with telemetry include metrology, data encoding, and data communications. Let's examine these steps in a communicating electric meter (often called a "smart meter"). The basic components of a smart meter include:

- An energy measurement element (metrology)
- A data processing element
- A communications processor

An electric meter

The energy measurement component in a meter takes a physical phenomenon and represents it as an electrical signal. Among these physical phenomena are the voltage and current of the electricity passing through the meter. The electrical signals in the metrology section of the meter are "analogs" of the physical phenomena, meaning electrical representations of physical quantities such as voltage and current. These measurements are usually calibrated against government standards (maintained in North America by Measurement Canada and the United States National Institute of Standards and Technology).

The next step is processing the electrical signals into digital data, which may be done by a microcontroller. Lastly, the digital data are communicated to an energy service provider for billing and optionally to devices in a home or building for energy monitoring and management. Communications may be performed by a dedicated processor inside the meter.

Figure 1 – Constituents of a Home or Building System Application

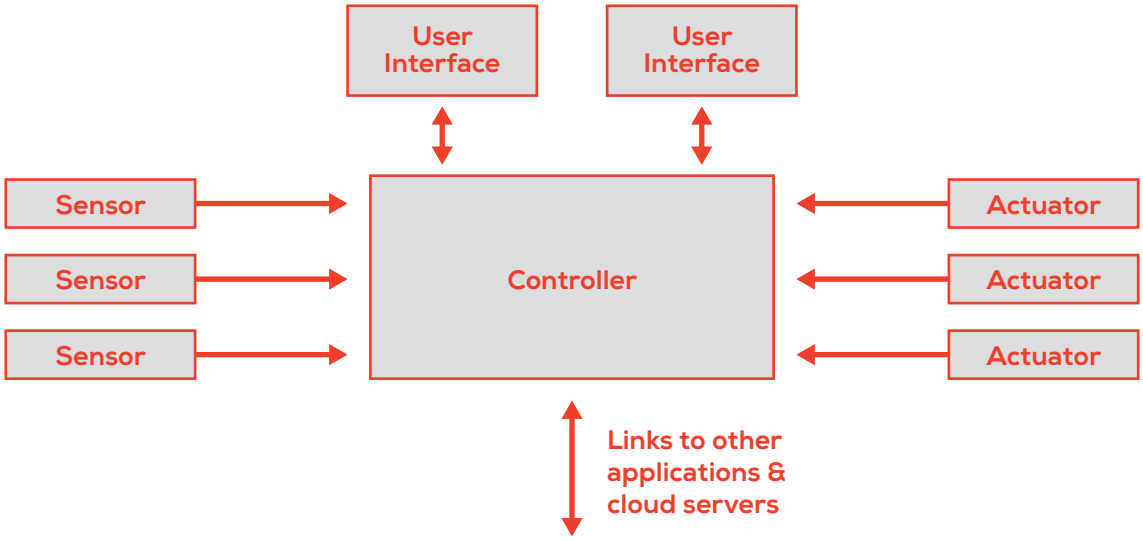

The three functions of measurement, data processing, and communications are found in most sensors. Low-cost sensors typically focus on measuring one physical parameter with specified accuracy and precision. In an inexpensive sensor, data encoding is minimal and communications often consists of an analog signal impressed on a wire to represent the measurement, rather than sending a digital data stream using a formal communications protocol.

A device intended to be connected to the Internet requires a more sophisticated communications interface than usually embedded in sensors and actuators. The phrase "Internet of Things" implies that devices such as sensors and actuators communicate via the Internet. However the design of the Internet poses challenges for such applications.

The goal of the Internet

The Internet was designed in the 1960s as a distributed network for sending text messages. This is similar to the telegraph network of the 1800s, but with a more robust and resilient infrastructure. Unlike the long-distance telephone network, which is based on a hierarchy of switching offices, the Internet uses distributed communications.

The design goal of the Internet was to deliver messages from computer terminals throughout the United States without errors, but with some delay, even if parts of the network were not functioning. The delay depends on network traffic and network performance, and might vary from message to message. This was acceptable for sending messages like telegrams where the messages were eventually read by a person. By the mid-70s, researchers were using computer terminals at major universities for exchanging messages with colleagues via the Internet (then called the ARPANET, an acronym for the U.S. government project that funded the research). This was the beginning of electronic mail, now called "e-mail."

A diversity of pathways is employed to overcome network failures. A message from a sender is divided into fixed-size packets for transmission through the Internet via routers: specialized communication processors. The Internet uses a mesh of routers acting as relay stations that receive packets from sender computers and from other routers, and pass these packets to other routers and to recipient computers. There are more that half a million routers located through the world.

The "Internet Protocol" or IP instructs routers to send packets in the direction of the recipient. Each packet is processed separately, may take a different route from sender to receiver, and may arrive out of the original sequence, as shown in Figure 2. Another process within the suite of Internet protocols (called the "Transport Control Protocol" or TCP) reassembles the packets received into the proper order in order to regenerate the message. TCP also checks for any damaged or missing

Figure 2 – Internet Transport of a Message as Packets
(Packets are re-ordered at the receiver)

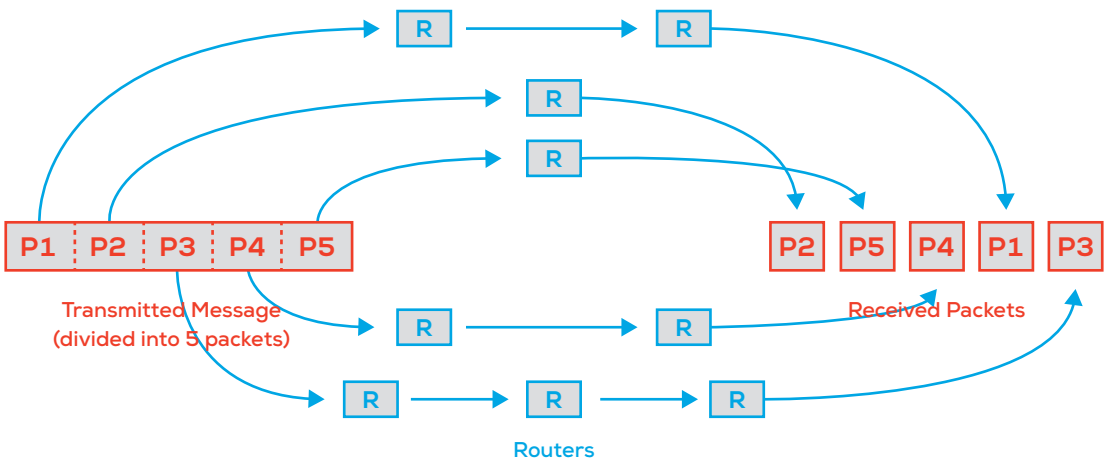

packets and corrects errors where possible or requests retransmission of these packets before delivering the message. All applications on the Internet use packet communications.

In the early 1990s, the World Wide Web was created using the messaging capabilities of the Internet. The World Wide Web (commonly called the "Web") is a remarkably useful tool for disseminating information especially in visual format between a data repository, called a server, and a user display. Messages intended for the Web include embedded codes (constituting a markup language) that instruct a computer display where on the screen to place the message, and what font, color, etc. to use. Also, images and graphics are encoded into messages for transmission through the Internet. Decoding software for image data is now built into computer operating systems and Web browsers.

Limitations of the Internet

Whether the Internet is carrying e-mail or Web data, the intended recipient is a person. We can easily tolerate some data delays in receiving e-mail or viewing Web pages. Machine-to-machine (M2M) communications, which is essential for the Internet of Things, was never a design goal of the Internet. M2M communications requires timely delivery of messages with minimum variation in message arrival times (called "jitter").

The challenge of extending the Internet to new real-time applications is evident in telephone calls via the Internet (called VoIP – Voice over Internet Protocol). VoIP calls are often choppy and distorted. Various workarounds to the limitations of the Internet communication protocols are being developed to improve VoIP.

The Internet is not the only protocol

Just as foreign diplomats engage with a host country using a formal protocol so they can understand each other, computer equipment from a variety of makers can exchange messages if they conform to an agreed protocol. A communications messaging protocol specifies the format of messages (called the "syntax"), the meaning of the message elements (called the "semantics"), and the sequence of messages including timing for delivery.

Messages are transmitted via unreliable media such as wires or radio using a communications media protocol. About 70 years ago, MIT Professor Claude Shannon[2] proved that reliable communications could be achieved over an unreliable medium if the data rate were kept below the "channel capacity." He left it for future engineers to develop protocols that could come close to achieving channel capacity with no errors. Thus, communications engineering addresses two fundamental issues:

- Source coding using message protocols
- Channel coding using media protocols

When multiple devices communicate via a common medium, a network is created. The media for the original

Internet were telephone lines. Because applications requiring communications are diverse and media performance varies considerably, there is no one optimum communications protocol. For example, tight timing specifications are essential for communications among devices and robots on a production line. For transferring computer files, time delays are not as important as accuracy. A few missing bytes can make a file unreadable.

So why is the world focusing on the Internet for such a diversity of communications ranging from Web browsing to VoIP to M2M? Because the Internet is deployed throughout the world and equipment to support the Internet is widely available at relatively low cost. However, no protocol is optimal for all applications. The Internet was designed for 1960s-era electronics. To adapt for the subsequent 50 years, the Internet has been patched with numerous clever workarounds. Some of these patches created and then attempted to remedy weakness that have exposed opportunities for cybersecurity attacks.

Communications for M2M

Eventually, telemetry applications for M2M communications will enable things to communicate with things to create a network of things. However, the Internet is not the optimal communications network for M2M.

Most home and building automation networks are not based on the Internet. They use local area networks developed for office and factory automation and adapted for the premises environment. Examples of specialized home and building networks are Echonet, KNX, LonTalk, and protocols specialized for a particular medium such as Wi-Fi and Z-Wave for radio or UPB for power lines.

Some Internet of Things devices will have embedded Internet protocols, but many will use *ad hoc* protocols for local networks. These devices may communicate remotely via the Internet using an external adapter (sometimes called a "proxy") or a gateway where the local messages are translated into the protocol of the Internet to reach a service provider server.

The re-branding of home automation

So the Internet of Things should more accurately be called the "Network of Things." But the Internet of Things has now become a marketing juggernaut, regardless of the underlying communications protocol.

The U.S. popular press started to run articles about home automation in the mid-80s with the prediction that "the year of home automation" would arrive by the late 80s. It has taken a little longer than expected! Coldwell Banker, the real estate brokerage firm, reported[3] in January 2016 that:

- Smart home technology is primed to become mainstream as soon as 2016.
- Smart home technology is no longer just for early adopters or the young and affluent.
- Homeowners are willing to invest in the technology if it will help their homes sell faster.

Coldwell Banker predicted that nearly 50 percent of Americans would have home automation technology by the end of 2016. Leading applications include entertainment, lighting, energy management, and safety.

Many exhibitors at the 2016 Consumer Electronics Show (CES) last January seemed to equate the Internet of Things with home automation. As Shakespeare said, "A rose by any other name would smell as sweet," which is interpreted by Wikipedia to mean that the names of things do not affect what they really are. So regardless of whether the Internet of Things is technically correct terminology, we should welcome these new entrants to the home automation industry. ●

Notes:

1. N. Gershenfeld, R. Krikorian, and D. Cohen, "The Internet of Things," *Scientific American*, 291:4, pp. 76–81, 2004.
2. Claude E. Shannon, "Communication in the Presence of Noise," Proceedings of the Institute of Radio Engineers, 37:1, January 1949, pp. 10–21.
3. http://blog.coldwellbanker.com/2016-is-the-year-smart-home-technology-will-be-mainstream

Dr. Kenneth Wacks has been a pioneer in establishing the home systems industry. He advises manufacturers and utilities worldwide on business opportunities, network alternatives, and product development in home and building systems. In 2008, the United States Department of Energy appointed him to the GridWise Architecture Council. For further information, please contact Dr. Wacks at 781.662.6211; kenn@alum.mit.edu; www.kenwacks.com.