1

Introduction

The Internet has been a great success over the past 20 years, growing from a small academic network into a global, ubiquitous network used regularly by over 1.4 billion people. It was the power of the Internet paradigm, tying heterogeneous networks together, and the innovative World Wide Web (WWW) model of uniform resource locators (URLs), the hypertext transfer protocol (HTTP) and universal content markup with the hypertext markup language (HTML) that made this possible. Grass-roots innovation has however been the most powerful driver behind the Internet success story. The Internet is open to innovation like no other telecommunication system before it. This has allowed all groups involved, from Internet architects to communication engineers, IT staff and everyday users to innovate, quickly adding new protocols, services and uses for Internet technology.

As the Internet of routers, servers and personal computers has been maturing, another Internet revolution has been going on – *The Internet of Things*. The vision behind the Internet of Things is that embedded devices, also called *smart objects*, are universally becoming IP enabled, and an integral part of the Internet. Examples of embedded devices and systems using IP today range from mobile phones, personal health devices and home automation, to industrial automation, smart metering and environmental monitoring systems. The scale of the Internet of Things is already estimated to be immense, with the potential of trillions of devices becoming IP-enabled. The impact of the Internet of Things will be significant, with the promise of better environmental monitoring, energy savings, smart grids, more efficient factories, better logistics, better healthcare and smart homes.

The Internet of Things revolution started in the 1990s with industrial automation systems. Early proprietary networks in industrial automation were quickly replaced by different forms of industrial Ethernet, and Internet protocols became widely used between embedded automation devices and back-end systems. This trend has continued in all other automation segments, with Ethernet and IP becoming ubiquitous. Machine-to-machine (M2M) telemetry made a breakthrough already in the early 2000s, with the use of cellular modems and IP to monitor and control a wide range of equipment from vending machines to water pumps. Building automation systems have gone from legacy control to making wide use of wired IP communications through the *Building Automation and Control Network* (BACnet) and

6LoWPAN: The Wireless Embedded Internet Zach Shelby and Carsten Bormann

Open Building Information Exchange (oBIX) standards. More recently, automatic metering infrastructures and smart grids are being deployed at a rapid rate, largely depending on the scalability and universal availability of IP technology. Finally, mobile phones have become almost universally IP-enabled embedded devices currently making up the largest body of devices belonging to the Internet of Things.

An equally important development has been happening in the services that are used to monitor and control embedded devices. Today these services are almost universally built on Internet technology, and more commonly are implemented using web-based services. *Web Service* technologies have completely changed the way business and enterprise applications are designed and deployed. It is this combination of Internet-connected embedded devices and Web-based services which makes the Internet of Things a powerful paradigm.

Hundreds of millions of embedded devices are already IP-enabled, but the Internet of Things is still in its infancy in 2009. Although the capabilities of processor, power and communications technology have continuously increased, so has the complexity of communications standards, protocols and services. Thus, so far, it has been possible to use Internet capabilities in only the most powerful embedded devices. Additionally, low-power wireless communications limits the practical bandwidth and duty-cycle available. Throughout the 1990s and early 2000s we have seen a large array of proprietary low-power embedded wireless radio and networking technologies. This has fragmented the market and slowed down the deployment of such technology.

The Institute of Electrical and Electronics Engineers (IEEE) released the 802.15.4 low-power wireless personal area network (WPAN) standard in 2003, which was a major milestone, providing the first global low-power radio standard. Soon after, the ZigBee Alliance developed a solution for ad hoc control networks over IEEE 802.15.4, and has produced a lot of publicity about the applications of wireless embedded technology. ZigBee and proprietary networking solutions that are vertically bound to a link-layer and application profiles only solve a small portion of the applications for wireless embedded networking. They also have problems with scalability, evolvability and Internet integration. A new paradigm was needed to enable low-power wireless devices with limited processing capabilities (see Figure 1.1) to participate in the Internet of Things, forming what we call the *Wireless Embedded Internet*.

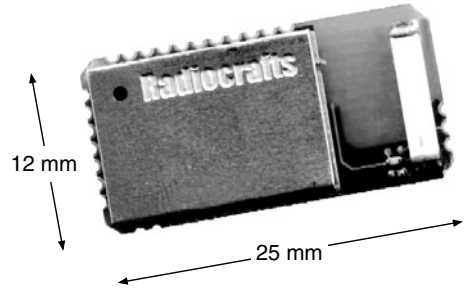

Figure 1.1 Wireless embedded 6LoWPAN device.

This book introduces a set of Internet standards which enable the use of *IPv6 over low-power wireless area networks* (6LoWPAN)[1], which is the key to realizing the Wireless Embedded Internet. 6LoWPAN breaks down the barriers to using IPv6 in low-power, processing-limited embedded devices over low-bandwidth wireless networks. IPv6, which is the newest version of the Internet Protocol, was developed in the late 1990s as a solution to the rapid growth and challenges facing the Internet. The further growth of the Internet of Things will be made possible thanks to IPv6.

In this chapter we give an overview of 6LoWPAN. First the Internet of Things is introduced, followed by the ideas behind 6LoWPAN, IETF standardization, related trends and applications of 6LoWPAN technology in Section 1.1. The overall 6LoWPAN architecture is then introduced in Section 1.2. A comprehensive overview of 6LoWPAN basic mechanisms and the link-layer are given in Section 1.3, followed by a 6LoWPAN network example in Section 1.4.

1.1 The Wireless Embedded Internet

What is the Internet of Things in practice? Maybe the simplest definition is that the Internet of Things encompasses all the embedded devices and networks that are natively IP-enabled and Internet-connected, along with the Internet services monitoring and controlling those devices. Figure 1.2 shows an illustration of the Internet of Things vision.

Today's Internet is made up of a *core Internet* of backbone routers and servers, including millions of nodes (any kind of network device) in total. The core Internet changes rarely and has extremely high capacity. The vast majority of today's Internet nodes are in what is sometimes called the *fringe Internet*. The fringe Internet includes all the personal computers, laptops and local network infrastructure connected to the Internet. This fringe changes rapidly, and is estimated to have up to a billion nodes. In 2008 it was estimated that the Internet had approximately 1.4 billion regular users, and Google announced that over a *trillion* unique URLs existed in their search indexes. The growth of the fringe is dependent on the number of Internet users and the personal devices used by them. The Internet of Things, sometimes referred to as the *embedded fringe*, is the biggest challenge and opportunity for the Internet today. It is made up of the IP-enabled embedded devices connected to the Internet, including sensors, machines, active positioning tags, radio-frequency identification (RFID) readers and building automation equipment to name but a few. The exact size of the Internet of Things is hard to estimate, as its growth is not dependent on human users. It is assumed that the Internet of Things will soon exceed the rest of the Internet in size (number of nodes) and will continue growing at a rapid rate. The long-term potential size of the Internet of Things is in trillions of devices. The greatest growth potential in the future comes from embedded, low-power, wireless devices and networks that until now have not been IP-enabled – the Wireless Embedded Internet. In 2008 the IP Smart Objects (IPSO) Alliance [IPSO] was formed by industry leaders to promote the use of Internet protocols by smart objects and the Internet of Things through marketing, education and interoperability.

The Wireless Embedded Internet is a subset of the Internet of Things, and the main subject of this book. We define the Wireless Embedded Internet to include resource-limited

[1] The 6LoWPAN acronym has been redefined on purpose in this book, as "Personal" is no longer relevant to the technology. WPAN originally referred to IEEE 802.15.4 Wireless Personal Area Network.

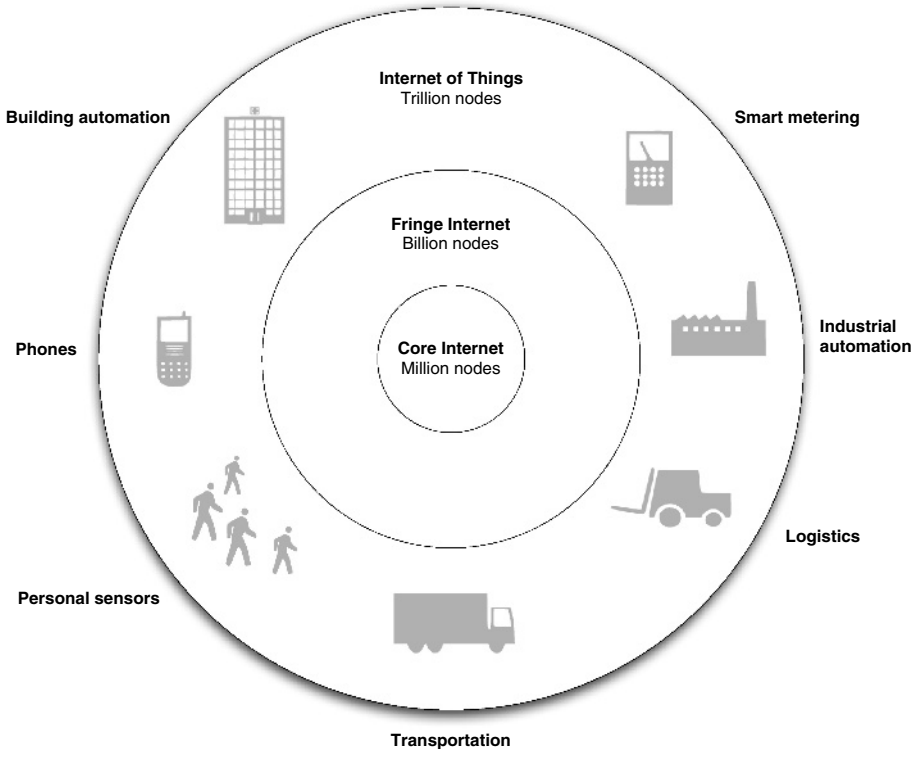

Figure 1.2 The Internet of Things vision.

embedded devices, often battery powered, connected by low-power, low-bandwidth wireless networks to the Internet. 6LoWPAN was developed to enable the Wireless Embedded Internet by simplifying IPv6 functionality, defining very compact header formats and taking the nature of wireless networks into account [6LoWPAN].

1.1.1 Why 6LoWPAN?

There are a huge range of applications which could benefit from a Wireless Embedded Internet approach. Today these applications are implemented using a wide range of proprietary technologies which are difficult to integrate into larger networks and with Internet-based services. The benefits of using Internet protocols in these applications, and thus integrating them with the Internet of Things include [RFC4919]:

- IP-based devices can be connected easily to other IP networks without the need for translation gateways or proxies.

- IP networks allow the use of existing network infrastructure.

- IP-based technologies have existed for decades, are very well known, and have been proven to work and scale. The *socket API* (application programming interface) is one of the most well-known and widely used APIs in the world.
- IP technology is specified in an open and free way, with standards processes and documents available to anyone. The result is that IP technology encourages innovation and is better understood by a wider audience.
- Tools for managing, commissioning and diagnosing IP-based networks already exist (although many management protocols need optimization for direct use with 6LoWPAN Nodes as we will discuss in Chapter 5).

Until now only powerful embedded devices and networks have been able to participate natively with the Internet. Direct communication with traditional IP networks requires many Internet protocols, often requiring an operating system to deal with the complexity and maintainability. Traditional Internet protocols are demanding for embedded devices for the following reasons:

Security: IPv6 includes optional support for IP Security (IPsec) [RFC4301] authentication and encryption, and web services typically make use of secure sockets or transport layer security mechanisms. These techniques may be too complex, especially for simple embedded devices.

Web services: Internet services today rely on web-services, mainly using the transmission control protocol (TCP), HTTP, SOAP and XML with complex transaction patterns.

Management: Management with the simple network management protocol (SNMP) and web-services is often inefficient and complex.

Frame size: Current Internet protocols require links with sufficient frame length (minimum of 1280 bytes for IPv6), and heavy application protocols require substantial bandwidth.

These requirements have in practice limited the Internet of Things to devices with a powerful processor, an operating system with a full TCP/IP stack, and an IP-capable communication link. Typical embedded Internet devices today include industrial devices with Ethernet interfaces, M2M gateways with cellular modems, and advanced smart phones. A large majority of embedded applications involve limited devices, with low-power wireless and wired network communications. Wireless embedded devices and networks are particularly challenging for Internet protocols:

Power and duty-cycle: Battery-powered wireless devices need to keep low *duty cycles* (the percentage of time active). The basic assumption of IP is that a device is always connected.

Multicast: Wireless embedded radio technologies, such as IEEE 802.15.4, do not typically support multicast, and flooding in such a network is wasteful of power and bandwidth. Multicast is crucial to the operation of many IPv6 features.

Mesh topologies: The applications of wireless embedded radio technology typically benefit from multihop mesh networking to achieve the required coverage and cost efficiency. Current IP routing solutions may not easily be applicable to such networks (discussed at length in Chapter 4).

Bandwidth and frame size: Low-power wireless embedded radio technology usually has limited bandwidth (on the order of 20–250 kbit/s) and frame size (on the order of 40–200 bytes). In mesh topologies, bandwidth further decreases as the channel is shared and is quickly reduced by multihop forwarding. The IEEE 802.15.4 standard has a 127-byte frame size, with layer-2 payload sizes as low as 72 bytes. The minimum frame size for standard IPv6 is 1280 bytes [RFC2460], thus requiring fragmentation.

Reliability: Standard Internet protocols are not optimized for low-power wireless networks. For example, TCP is not able to distinguish between packets dropped because of congestion or packets lost on wireless links. Further unreliability occurs in wireless embedded networks because of node failure, energy exhaustion and sleep duty cycles.

The IETF 6LoWPAN working group [6LoWPAN] was created to tackle these problems, and to specifically enable *IPv6* to be used with wireless embedded devices and networks. Features of the IPv6 design such as a simple header structure, and its hierarchical addressing model, made it ideal for use in wireless embedded networks with 6LoWPAN. Additionally, by creating a dedicated group of standards for these networks, the minimum requirements for implementing a lightweight IPv6 stack with 6LoWPAN could be aligned with the most minimal devices. Finally by designing a version of Neighbor Discovery (ND) specifically for 6LoWPAN, the particular characteristics of low-power wireless mesh networks could be taken into account. The result of 6LoWPAN is the efficient extension of IPv6 into the wireless embedded domain, thus enabling *end-to-end* IP networking and features for a wide range of embedded applications. Refer to [RFC4919] for the detailed assumptions, problem statement and goals of early 6LoWPAN standardization. Although 6LoWPAN was targeted originally at IEEE 802.15.4 radio standards and assumed layer-2 mesh forwarding [RFC4944], it was later generalized for all similar link technologies, with additional support for IP routing in [ID-6lowpan-hc, ID-6lowpan-nd].

1.1.2 6LoWPAN history and standardization

6LoWPAN is a set of standards defined by the Internet Engineering Task Force (IETF), which creates and maintains all core Internet standards and architecture work. A straightforward technical definition of 6LoWPAN would be:

> 6LoWPAN standards enable the efficient use of IPv6 over low-power, low-rate wireless networks on simple embedded devices through an adaptation layer and the optimization of related protocols.

The IETF 6LoWPAN working group was officially started in 2005, although the history of embedded IP goes back farther. Throughout the 1990s it was assumed that Moore's law would advance computing and communication capabilities so rapidly that soon any embedded device could implement IP protocols. Although partially true, and the Internet of Things has grown rapidly, it did not hold for cheap, low-power microcontrollers and low-power wireless radio technologies. The vast majority of simple embedded devices still make use of 8-bit and 16-bit microcontrollers with very limited memory, as they are low-power, small and cheap. At the same time, the physical trade-offs of wireless technology have resulted in short-range, low-power wireless radios which have limited data rates, frame sizes and duty cycles such as

in the IEEE 802.15.4 standard. Early work on minimizing Internet protocols for use with low-power microcontrollers and wireless technologies includes μIP from the Swedish Institute of Computer Science [Dunkels03] and NanoIP from the Centre for Wireless Communications [Shel03]. The IEEE 802.15.4 standard released in 2003 was the biggest factor leading to 6LoWPAN standardization. For the first time a global, widely supported standard for low-power wireless embedded communications was available [IEEE802.15.4]. The popularity of this new standard gave the Internet community the needed encouragement to standardize an IP adaptation for such wireless embedded links.

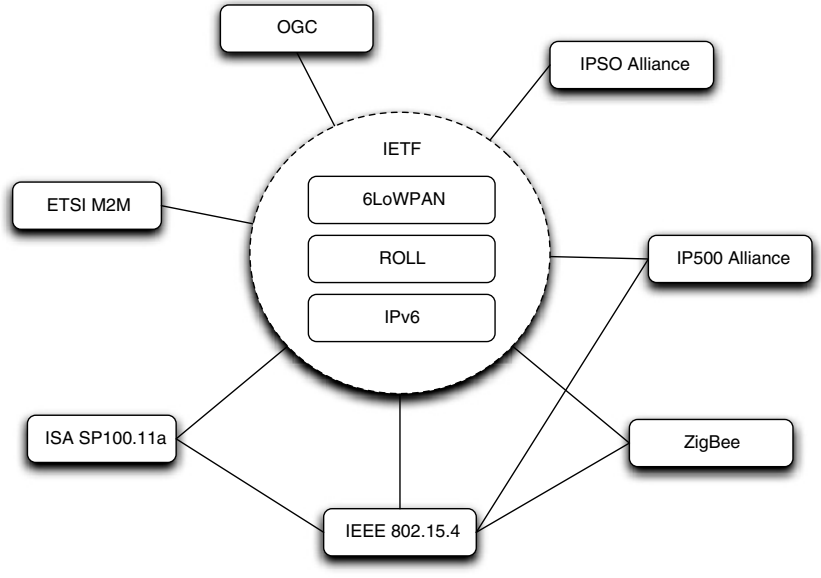

Figure 1.3 The relation of 6LoWPAN to related standards and alliances.

The first 6LoWPAN specifications were released in 2007, first with an informational RFC [RFC4919] specifying the underlying requirements and goals of the initial standardization, and then with a standard track RFC [RFC4944] specifying the 6LoWPAN format and functionality. Through experience with implementations and deployments, the 6LoWPAN working group continued with improvements to header compression [ID-6lowpan-hc], 6LoWPAN Neighbor Discovery [ID-6lowpan-nd], use cases [ID-6lowpan-uc] and routing requirements [ID-6lowpan-rr]. In 2008 a new IETF working group was formed, *Routing over Low-power and Lossy Networks* (ROLL)[ROLL]. This working group specifies routing requirements and solutions for low-power, wireless, unreliable networks. Although not restricted to use with 6LoWPAN, that is one main target.

In 2008 ISA began standardization of a wireless industrial automation system called SP100.11a (also known as ISA100), which is based on 6LoWPAN. An overview of ISA100 is given in Chapter 7. Recent activities related to 6LoWPAN include the IP for Smart Objects (IPSO) Alliance founded in 2008 to promote the use of IP in smart objects and Internet

of Things business [IPSO], and the IP500 Alliance which is developing a recommendation for 6LoWPAN over IEEE 802.15.4 sub-GHz radio communications [IP500]. Figure 1.3 shows the relations between related standards bodies and alliances. The Open Geospatial Consortium (OGC) specifies IP-based solutions for geospatial and sensing applications. In 2009 the European Telecommunication Standards Institute (ETSI) [ETSI] started a working group for standardizing M2M, which includes an end-to-end IP architecture compatible with 6LoWPAN.

1.1.3 Relation of 6LoWPAN to other trends

There are several other trends to take into consideration when thinking about the Internet of Things. These include ZigBee, machine-to-machine (M2M) communications, the Future Internet, and wireless sensor networks (WSNs). This section looks into how each of these trends relates to the Internet of Things and 6LoWPAN in particular.

ZigBee is a protocol specification from an industry special interest group called the ZigBee Alliance, specializing in ad hoc control [ZigBee]. ZigBee was started in 2003 in conjunction with IEEE 802.15.4 standardization [IEEE802.15.4], and specifies a vertical protocol stack solution with similarities to Bluetooth. The protocol mainly makes use of IEEE 802.15.4 features, adding ad hoc networking, service discovery and application protocol profiles on top of that. ZigBee has been successful for multi-vendor ad hoc applications such as home automation. The ZigBee approach does have several down-sides, including reliance on a single wireless link technology, tight coupling with application profiles, along with Internet integration and scalability limitations. In 2009 the ZigBee Alliance announced that ZigBee will start to integrate IETF standards such as 6LoWPAN and ROLL into its future specifications. Earlier work has shown how ZigBee application profiles can be carried over UDP/IP and 6LoWPAN [ID-tolle-cap], which is covered in more detail in Section 5.4.3. The integration of IP technology into ZigBee provides a much wider range of networking possibilities, beyond just ad hoc control.

Machine-to-machine (M2M) communications has become a popular industry term for the remote monitoring and control of machines over the Internet. Traditionally, M2M systems include M2M modules (usually a cellular modem) integrated into embedded devices together with an Internet-based back-end system. The M2M module measures and controls the device, and communicates over IP with the back-end M2M service. More recently, M2M gateways to local embedded networked devices have become more common. Thanks to native IP, 6LoWPAN networks can be connected to M2M services through simple routers and thus 6LoWPAN can be considered to be a natural extension of M2M. Machine-to-machine communications has been an important driving force in the development and growth of the Internet of Things, which has continued with the ETSI M2M standardization effort.

The Future Internet [Bauge08] is a term used to describe research into what the Internet architecture and protocols could look like in 10–20 years. The US National Science Foundation has a long-term initiative on Future Internet Design (FIND) which covers network architecture, principles as well as mechanism design [FIND]. Several European projects specialize in Future Internet research, for example the EU 4WARD project [4WARD], in cooperation with the European Future Internet Assembly [FIAssembly]. Although most of the research related to Future Internet does not consider embedded devices and networks, this aspect is starting to gain interest. The EU SENSEI project [SENSEI] for example specializes

in making wireless sensor and embedded networks a part of the global Internet, both current and future. One of the subjects of the project is how wireless embedded networks and 6LoWPAN type functionality can be made an integral part of the Future Internet. Several examples throughout this book are taken from the SENSEI project as it has been doing leading work in this area.

The term Wireless Sensor Network (WSN) comes from an academic movement starting in the mid 1990s into research on low-power ad hoc wireless networked sensors and actuators. The US government was very interested in the application of low-power sensing in military and security applications, and provided extensive funding for the subject. The research area later developed into a widely popular subject with a large range of applications, and a huge collection of results and trials. These networks have traditionally been thought to be completely isolated, and thus typically Internet compatibility or standards were not taken into consideration. Instead each project has tended to produce its own optimized wireless, network and algorithm solutions. Additionally most of the envisioned applications of sensor networks were created by university researchers, and they most often did not have a real market need. More recently the importance of standards, marketable applications and the importance of Internet services have encouraged the WSN community to become involved with 6LoWPAN standardization and the IPSO Alliance. The result is that a lot of the innovation produced through WSN research is starting to be applied to Wireless Embedded Internet technology, a good example being the IETF ROLL working group.

There is a strong trend of convergence in standardization, industry and research, as indicated above. This convergence is clearly steering towards an Internet-based approach as the requirements of modern-day embedded applications clearly demand it. 6LoWPAN has been a result of and catalyst for convergence to the Internet of Things.

1.1.4 Applications of 6LoWPAN

The reason why there are such a large number of technical solutions in the wireless embedded networking market is that the requirements, scale and market of embedded applications vary wildly. Applications can range from personal health sensor monitoring to large scale facility monitoring, which differ greatly. This is in contrast to PC information technology, which is fairly homogeneous and mainly aimed at home and office environments. The ideal use of 6LoWPAN is in applications where:

- embedded devices need to communicate with Internet-based services,
- low-power heterogeneous networks need to be tied together,
- the network needs to be open, reusable and evolvable for new uses and services, and
- scalability is needed across large network infrastructures with mobility.

Connecting the Internet to the physical world enables a wide range of interesting applications where 6LoWPAN technology may be applicable, for example:

- home and building automation
- healthcare automation and logistics

- personal health and fitness (see Figure 1.4)
- improved energy efficiency
- industrial automation (see Figure 1.5)
- smart metering and smart grid infrastructures
- real-time environmental monitoring and forecasting
- better security systems and less harmful defense systems
- more flexible RFID infrastructures and uses
- asset management and logistics
- vehicular automation

Figure 1.4 Example of a personal fitness monitoring application. (Reproduced by Permission of © SENSEI Consortium.)

One interesting example application of 6LoWPAN is in *facility management*, which is the management of large facilities using a combination of building automation, asset management and other embedded systems. This quickly growing field can benefit from 6LoWPAN, is feasible with today's technology, and has real business demand. For these reasons it is an ideal example, which is introduced in the next section.

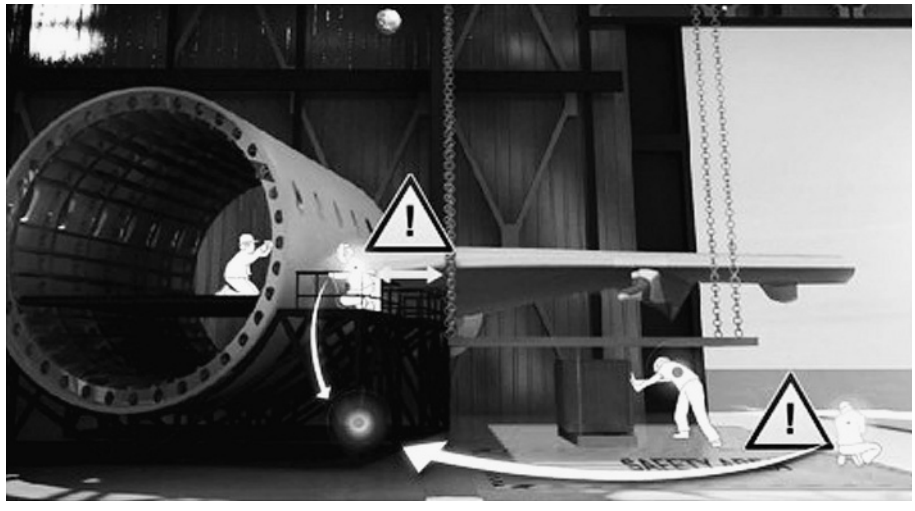

Figure 1.5 Example of an industrial safety application. (Reproduced by Permission of © SENSEI Consortium.)

1.1.5 Example: facility management

Facility management is a very interesting application for the Internet of Things, and is one use case that has been examined in detail by the SENSEI project [SENSEI]. It involves the integrated management of building facilities. Facility management services are becoming more common, and are typically web-based. Figure 1.6 shows a facility management use case from the SENSEI project. Wireless embedded networking has a large range of applications in facility management including:

Door access control: Access control involves the use of RFID or active tag based identifiers to control and log the access to different parts of a building automatically.

Building automation: Building automation involves the use of sensors and control to improve the operations and efficiency of a building.

Tracking: Tracking involves the use of active tags on people, equipment and supplies which are tracked by the wireless infrastructure throughout a facility. Tracking results are used in asset management, security and logistics optimization.

Energy reduction: Energy reduction in facilities can be achieved through intelligent lighting control, heating control, ventilation and air conditioning control, and the automatic power control of electric equipment.

Maintenance: The maintainability of facilities can be improved through the remote monitoring of the building itself and the systems in the building which today are typically monitored manually.

Smart metering: The use of resources in large facilities can be reduced and better controlled through more intelligent metering of electricity, gas and water using an *automatic metering infrastructure* (AMI).

The stakeholders in facility management include the providers of intelligent facility management systems and services, users of these services and third parties. The providers of facility management services play an important role as a huge amount of data needs to be collected, processed and leveraged to provide the services required in a beneficial way. The automation systems in facilities may include access control, building automation, tracking, maintenance monitoring and metering systems. Users of facility management include building owners or renters, building users and facility managers. Additionally many third parties are involved with facility management such as security companies, insurance companies and utilities. Some of these stakeholders are identified in Figure 1.6.

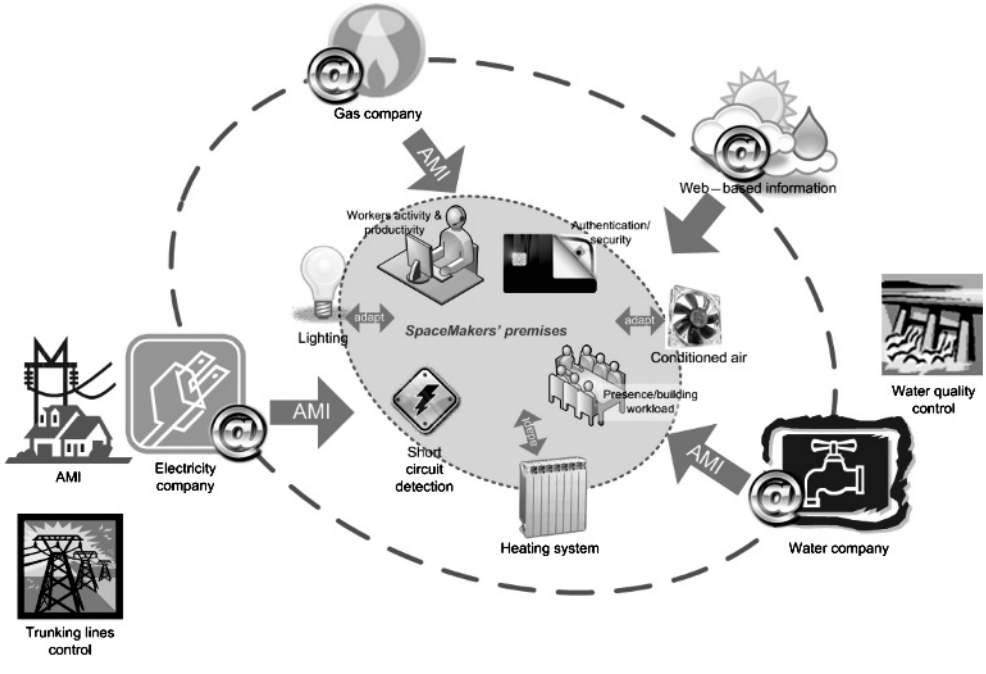

Figure 1.6 An example of a facility management system including an automatic metering infrastructure (AMI). (Reproduced by Permission of © SENSEI Consortium.)

The main rationales for facility management are improvements in energy and resource efficiency, an increase in worker productivity, and more secure and comfortable buildings. Buildings are major consumers of energy: it is estimated that in the EU and the USA, 40 percent of all energy is consumed in the building sector [Baden06, DoE06], and that carbon emissions could be reduced by 22 percent through improved efficiency [2002/91/EC]. For the enterprise users of buildings, an even more important benefit is improved worker efficiency along with better comfort and security in general. Substantial cost savings may be possible through productivity improvements.

Facility management provides many technical challenges for embedded devices and networking. The large range of systems to be integrated needs interoperability between systems, as well as network integration of heterogeneous technology. Furthermore new devices and applications will be added over time, so evolvability is important. The scalability of wireless embedded networking in large buildings is demanding. The density of devices in a single space can reach hundreds of nodes, and there is a mix of fixed and mobile devices across a large area. Battery-powered wireless devices require intelligent networking designed to maximize the lifetime of devices, and thus reduce maintenance. Facility management systems and devices must be cost-efficient, and installation straightforward compared to the long-term benefits achieved through these services. Finally, although privacy is easier in enterprise networks, security is a challenging aspect when applying wireless embedded networking. We will consider how to apply 6LoWPAN to solve networking requirements such as these throughout this book.

1.2 The 6LoWPAN Architecture

The Wireless Embedded Internet is created by connecting islands of wireless embedded devices, each island being a *stub network* on the Internet. A stub network is a network which IP packets are sent from or destined to, but which doesn't act as a transit to other networks. The 6LoWPAN architecture is made up of *low-power wireless area networks* (LoWPANs)[2], which are IPv6 stub networks. The overall 6LoWPAN architecture is presented in Figure 1.7. Three different kinds of LoWPANs have been defined: Simple LoWPANs, Extended LoWPANs, and Ad hoc LoWPANs. A LoWPAN is the collection of 6LoWPAN Nodes which share a common IPv6 address *prefix* (the first 64 bits of an IPv6 address), meaning that regardless of where a node is in a LoWPAN its IPv6 address remains the same. An *Ad hoc LoWPAN* is not connected to the Internet, but instead operates without an infrastructure. A *Simple LoWPAN* is connected through one *LoWPAN Edge Router* to another IP network. A *backhaul link* (point-to-point, e.g. GPRS) is shown in the figure, but this could also be a *backbone link* (shared). An *Extended LoWPAN* encompasses the LoWPANs of multiple edge routers along with a backbone link (e.g. Ethernet) interconnecting them.

LoWPANs are connected to other IP networks through *edge routers*, as seen in Figure 1.7. The edge router plays an important role as it routes traffic in and out of the LoWPAN, while handling 6LoWPAN compression and Neighbor Discovery for the LoWPAN. If the LoWPAN is to be connected to an IPv4 network, the edge router will also handle IPv4 interconnectivity (discussed further in Section 4.3). Edge routers typically have management features tied into overall IT management solutions. Multiple edge routers can be supported in the same LoWPAN if they share a common backbone link.

A LoWPAN consists of nodes, which may play the role of host or router, along with one or more edge routers. The network interfaces of the nodes in a LoWPAN share the same IPv6 prefix which is distributed by the edge router and routers throughout the LoWPAN. In order to facilitate efficient network operation, nodes register with an edge router. These operations are part of *Neighbor Discovery* (ND), which is an important basic mechanism

[2]The terms 6LoWPAN and LoWPAN are often used interchangeably. In this book we use 6LoWPAN as a general term for the technology or set of standards, and the term LoWPAN as it is used in the IETF standards: to refer to a specific type of LoWPAN or node.

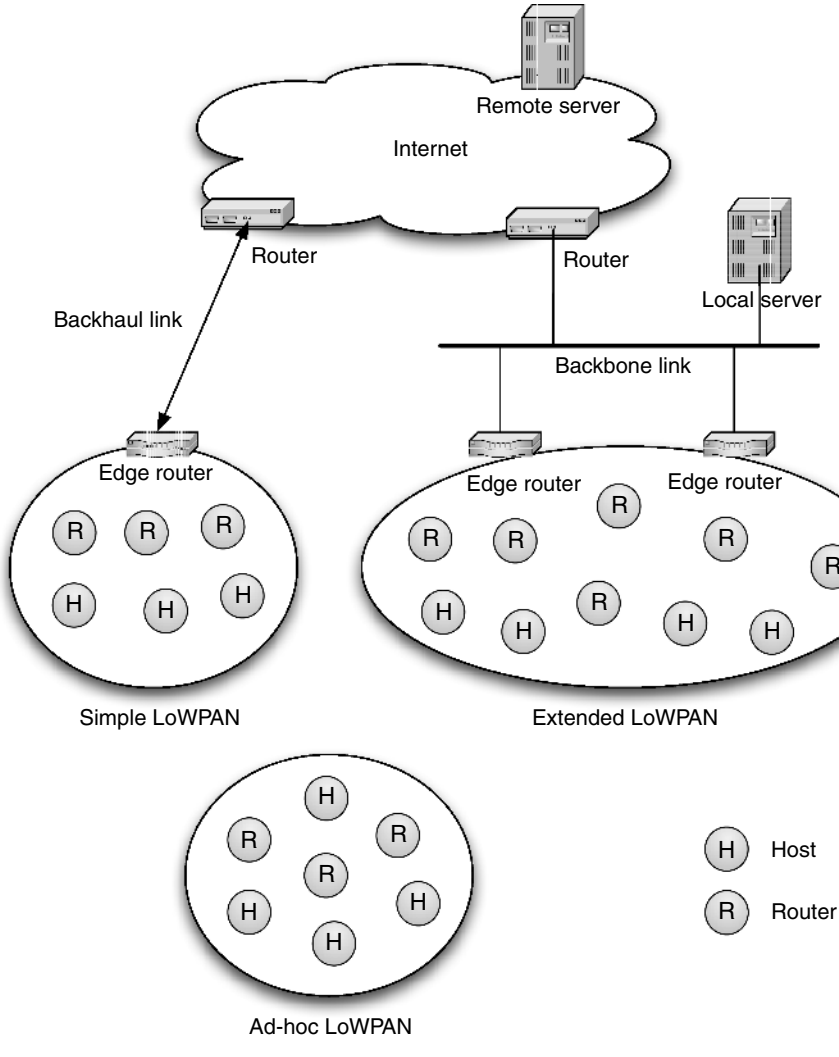

Figure 1.7 The 6LoWPAN architecture.

of IPv6. Neighbor Discovery defines how hosts and routers interact with each other on the same link. LoWPAN Nodes may participate in more than one LoWPAN at the same time (called *multi-homing*), and fault tolerance can be achieved between edge routers. LoWPAN Nodes are free to move throughout the LoWPAN, between edge routers, and even between LoWPANs. Topology change may also be caused by wireless channel conditions, without physical movement. A multihop mesh topology within the LoWPAN is achieved either through link-layer forwarding (called *Mesh-Under*) or using IP routing (called *Route-Over*). Both techniques are supported by 6LoWPAN.

Communication between LoWPAN Nodes and IP nodes in other networks happens in an end-to-end manner, just as between any normal IP nodes. Each LoWPAN Node is identified

by a unique IPv6 address, and is capable of sending and receiving IPv6 packets. Typically LoWPAN Nodes support ICMPv6 traffic such as “ping”, and use the user datagram protocol (UDP) as a transport. In Figure 1.7 the Simple LoWPAN and Extended LoWPAN Nodes can communicate with either of the servers through their edge router. As the payload and processing capabilities of LoWPAN Nodes are extremely limited, application protocols are usually designed using a simple binary format in a UDP payload. Application protocols suitable for 6LoWPAN are discussed in Chapter 5.

The main difference between a Simple LoWPAN and an Extended LoWPAN is the existence of multiple edge routers in the LoWPAN, which share the same IPv6 prefix and a common backbone link. Multiple LoWPANs can overlap each other (even on the same channel). When moving from one LoWPAN to another, a node’s IPv6 address will change. A LoWPAN Edge Router is typically connected to the Internet over a backhaul link such as cellular or DSL [ID-6lowpan-nd]. A network deployment may also choose to use multiple Simple LoWPANs rather than an Extended LoWPAN on a shared backbone link, e.g. for management reasons. This is not a problem if there is low mobility between LoWPANs in the network, or the application does not assume stable IPv6 addresses for nodes. A deployment example of a Simple LoWPAN connected by a backhaul link to the Internet is given in Section 1.4.

In an Extended LoWPAN configuration, as shown on the right-hand side of Figure 1.7, multiple edge routers share a common backbone link and collaborate by sharing the same IPv6 prefix, offloading most Neighbor Discovery messaging to the backbone link [ID-6lowpan-nd]. This greatly simplifies LoWPAN Node operation as IPv6 addresses are stable throughout the Extended LoWPAN and movement between edge routers is very simple. Edge routers also handle IPv6 forwarding on behalf of the nodes. To IP nodes outside the LoWPAN, the LoWPAN Nodes are always reachable regardless of their attachment point in the Extended LoWPAN. This enables large enterprise 6LoWPAN infrastructures to be built, functioning similar to a WLAN (WiFi) access point infrastructure (but at layer 3 instead of layer 2).

6LoWPAN does not require an infrastructure to operate, but may also operate as an Ad hoc LoWPAN [ID-6lowpan-nd]. In this topology, one router must be configured to act as a simplified edge router, implementing two basic functionalities: unique local unicast address (ULA) generation [RFC4193] and handling 6LoWPAN Neighbor Discovery registration functionality. From the LoWPAN Node point of view the network operates just like a Simple LoWPAN, except the prefix advertised is an IPv6 local prefix rather than a global one, and there are no routes outside the LoWPAN.

LoWPAN types and 6LoWPAN Neighbor Discovery operation are covered in detail in Chapter 3. Also refer to the 6LoWPAN Neighbor Discovery document in [ID-6lowpan-nd] for the complete specification.

1.3 6LoWPAN Introduction

This section gives a short but comprehensive introduction to the core 6LoWPAN subjects covered in this book. The protocol stack, link-layer technology, addressing and header format are first explained, followed by bootstrapping, mesh topologies, and Internet integration.

1.3.1 The protocol stack

Figure 1.8 shows the IPv6 protocol stack with 6LoWPAN in comparison with a typical IP protocol stack and the corresponding five layers of the *Internet Model* (the four-layer model of [RFC1122] with a physical layer separated out of the link layer). The Internet Model is sometimes referred to as a "narrow waist" model, as the Internet Protocol ties together a wide variety of *link-layer* technologies with multiple transport and application protocols. A simple IPv6 protocol stack with 6LoWPAN (also called a 6LoWPAN protocol stack) is almost identical to a normal IP stack with the following differences. First of all 6LoWPAN only supports IPv6, for which a small adaptation layer (called the *LoWPAN adaptation layer*) has been defined to optimize IPv6 over IEEE 802.15.4 and similar link layers in [RFC4944]. In practice, 6LoWPAN stack implementations in embedded devices often implement the LoWPAN adaptation layer together with IPv6, thus they can alternatively be shown together as part of the network layer (see Section 6.2 for more about stack implementation issues). The most common transport protocol used with 6LoWPAN is the user datagram protocol (UDP) [RFC0768], which can also be compressed using the LoWPAN format. The transmission control protocol (TCP) is not commonly used with 6LoWPAN for performance, efficiency and complexity reasons. The Internet control message protocol v6 (ICMPv6) [RFC4443] is used for control messaging, for example ICMP echo, ICMP destination unreachable and Neighbor Discovery messages. Application protocols are often application specific and in binary format, although more standard application protocols are becoming available. Application protocols are discussed in detail in Chapter 5.

Adaptation between full IPv6 and the LoWPAN format (described later in this section) is performed by routers at the edge of 6LoWPAN islands, referred to as edge routers. This transformation is transparent, efficient and stateless in both directions. LoWPAN adaptation in an edge router typically is performed as part of the 6LoWPAN network interface driver and is usually transparent to the IPv6 protocol stack itself. Figure 1.9 illustrates one realization of an edge router with 6LoWPAN support. See Section 6.4 for edge router implementation considerations. Inside the LoWPAN, hosts and routers do not actually need to work with full IPv6 or UDP header formats at any point as all compressed fields are implicitly known by each node.

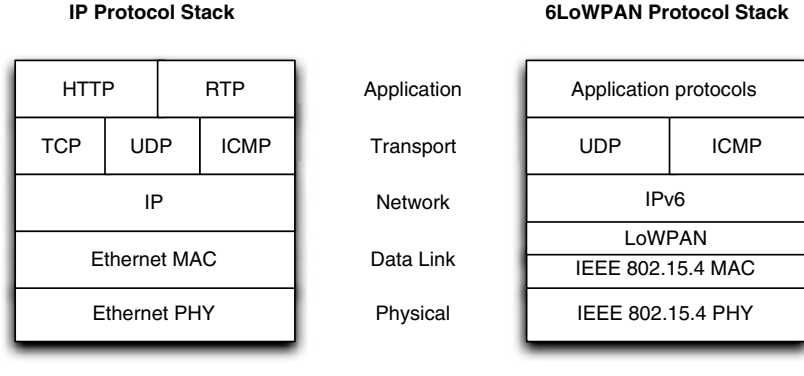

Figure 1.8 IP and 6LoWPAN protocol stacks.

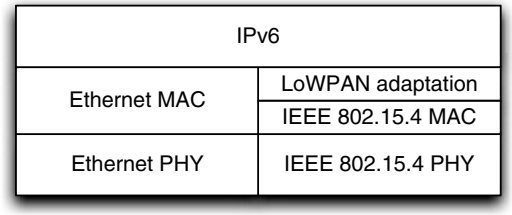

Figure 1.9 IPv6 edge router with 6LoWPAN support.

1.3.2 Link layers for 6LoWPAN

One of the most important functions of the Internet Protocol is the interconnection of heterogeneous links into a single interoperable network, providing a universal "narrow waist". This is equally true for 6LoWPAN and embedded networks, where there are many wireless (and also wired) link-layer technologies in use. The specialized applications of embedded networks require a wider range of communication solutions than typical personal computer networks, which almost universally use Ethernet and WiFi. Luckily the IEEE 802.15.4 standard is the most common 2.4 GHz wireless technology for embedded networking applications, and has been used as a baseline for 6LoWPAN development. Other technologies used with 6LoWPAN include sub-GHz radios, long-range telemetry links and even power-line communications. The requirements and interactions of 6LoWPAN with the link layer are discussed next, along with an introduction to IEEE 802.15.4, a sub-GHz radio and power-line communications.

There is a set of required or recommended features that a link should provide in order to work with Internet protocols. These include framing, addressing, error checking, length indication, some reliability, broadcast and a reasonable frame size. The issues involved with designing a subnetwork for use with IP are discussed in [RFC3819]. 6LoWPAN is designed to be used with a special type of link, and has its own set of link requirements and recommendations.

The most basic requirements for a link layer to support 6LoWPAN are framing, unicast transmission and addressing. Addressing is required to differentiate between nodes on a link, and to form IPv6 addresses which are then elided by 6LoWPAN compression. It is highly recommended that a link supports unique addresses by default (e.g. a 64-bit *extended unique identifier* [EUI-64]), to allow for stateless autoconfiguration. Multi-access links should provide a broadcast service. Multicast service is required by standard IPv6, but not by 6LoWPAN (broadcast is sufficient). IPv6 requires a *maximum transmission unit* (MTU) of 1280 bytes from a link, which 6LoWPAN fulfills by supporting fragmentation at the LoWPAN adaptation layer. A link should provide payload sizes at least 30 bytes in length to be useful (and preferably larger than 60 bytes). Although UDP and ICMP include a simple 16-bit checksum, it is recommended that the link layer also provides strong error checking. Finally, as IPsec may not always be practical for 6LoWPAN, it is highly recommended that links include strong encryption and authentication. The 2006 version of the IEEE 802.15.4

standard actually does not include a "next protocol identifier", making the detection of which protocol is being carried difficult. Although partially dealt with in the LoWPAN format using a *dispatch value*, it is a feature that a link should preferably have. Subnetwork design and link-layer issues are discussed in Section 2.2.

The next sections introduce three link-layer technologies used with 6LoWPAN: IEEE 802.15.4, a sub-GHz ISM band radio and low-rate power line communications.

IEEE 802.15.4

The IEEE 802.15.4 standard [IEEE802.15.4] defines low-power wireless embedded radio communications at 2.4 GHz, 915 MHz and 868 MHz. The first version of the standard was released in 2003, and was then revised in 2006. More recently the IEEE 802.15.4a standard was released, extending 802.15.4 with two new physical layer options: *Chirp Spread Spectrum* at 2.4 GHz and *Ultra Wide-Band* at 3.1–10.6 GHz. Work continues on new features such as MAC improvements in IEEE 802.15.4 Task Group 4e (TG4e), active RFID (TG4f), larger networks (TG4a) and specialized PHYs for China (TG4c) and Japan (TG4d). More information is available on these efforts from [IEEE]. In practice IEEE 802.15.4 at 2.4 GHz is used almost exclusively today as it provides reasonable data rates, and can be used globally. The sub-GHz channels are limited geographically with 915 MHz mainly available in North America and 868 MHz in the European Union (EU). That, combined with the limited data rates and channel selection of sub-GHz IEEE 802.15.4, means that there are only a few chips on the market today. Often more flexible sub-GHz chips tend to be used, as explained in the next section. This trend may yet change, with new sub-GHz applications becoming widespread and efforts like the IP500 Alliance, together with improvements in the latest IEEE 802.15.4 standard for sub-GHz channels.

The 802.15.4 standard provides 20–250 kbit/s data rates depending on the frequency. Channel sharing is achieved using carrier sense multiple access (CSMA), and acknowledgments are provided for reliability. Link-layer security is provided with 128-bit AES encryption. Addressing modes for 64-bit (long) and 16-bit (short) addresses are provided with unicast and broadcast capabilities. The physical layer payload is up to 127 bytes, with 72–116 bytes of payload available after link-layer framing, addressing, and optional security. The MAC can be run in two modes: beaconless mode and beacon-enabled mode. Beaconless mode uses pure CSMA channel access and operates quite like IEEE 802.11 without channel reservations. Beacon-enabled mode uses a hybrid time division multiple access (TDMA) approach, with the possibility of reserving time-slots for critical data. IEEE 802.15.4 includes many mechanisms for forming networks, and for controlling the superframe settings. An IEEE 802.15.4 reference is provided in Appendix B.

Early 6LoWPAN standardization work was originally aimed at the IEEE 802.15.4 standard [RFC4919, RFC4944] and thus assumed that some 802.15.4-specific features such as beacon-enabled mode and association mechanisms would be used along with 802.15.4 device roles. Based on practical experience with [RFC4944] and industry needs, recent 6LoWPAN standardization has been generalized to work with a larger range of link layers and avoids the assumption of IEEE 802.15.4-specific features. The use of 6LoWPAN with IEEE 802.15.4 is covered in more detail in Section 2.2.

Sub-GHz ISM band radios

Sub-GHz radio technologies using the industrial, scientific and medical (ISM) bands for unlicensed operation are especially popular in low-power wireless embedded applications such as telemetry, metering and remote control. The sub-GHz ISM bands cover 433 MHz, 868 MHz and 915 MHz. The main reasons for sub-GHz popularity are the better penetration of lower frequency, resulting in better range compared to 2.4 GHz, and the 2.4 GHz ISM band becoming very crowded in urban environments. One example of a popular sub-GHz chip is the Texas Instruments CC1101 transceiver [CC1101]. This transceiver acts as a reconfigurable radio and is capable of 300–928 MHz operation, with a wide variety of modulations, channel and data rates up to 500 kbit/s. Such a chip can also be used with an external power amplifier for increasing range. The features of the chip include carrier sensing, received signal strength indicator (RSSI) support, and frame sizes up to 250 bytes. The system-on-a-chip version, the CC1110, additionally includes a 128-bit AES encryption hardware engine. This kind of transceiver only provides the physical layer, so the data-link layer is implementation specific and needs to provide e.g. framing, addressing, error checking, acknowledgments and frame length. When designing a link layer for this type of transceiver, the IEEE 802.15.4 frame structure and beaconless mode operation is typically used as a starting point.

Power line communications

6LoWPAN also has interesting uses over special wired communication links, such as low-rate *power line communications* (PLC). Applications of this technology include home automation, energy efficiency monitoring and smart metering. One such system from Watteco [Watteco] uses what is called a watt pulse communication (WPC) technique, greatly reducing the complexity of communications. The data rate of the physical layer provided using WPC is 9.6 kbit/s, and the resulting channel over the power system of a house, building or urban area is multi-access and similar to a wireless CSMA channel. Watteco provides a version of WPC with an emulation of the IEEE 802.15.4 data link layer. This allows 6LoWPAN to be used with PLC in a very similar way to IEEE 802.15.4 and other ISM band radios. With PLC, multihop routing is not an issue as typically all nodes are on the same stable link. Multihop forwarding may be useful to interconnect several PLC subnets, or to integrate PLC and wireless 6LoWPAN islands.

1.3.3 Addressing

IP addressing with 6LoWPAN works just like in any IPv6 network, and is similar to addressing over Ethernet networks as defined by [RFC2464]. IPv6 addresses are typically formed automatically from the prefix of the LoWPAN and the link-layer address of the wireless interfaces. The difference in a LoWPAN is with the way low-power wireless technologies support link-layer addressing; a direct mapping between the link-layer address and the IPv6 address is used for achieving compression. This will be explained in Section 1.3.4.

Low-power wireless radio links typically make use of flat link-layer addressing for all devices, and support both unique long addresses (e.g. EUI-64) and configurable short addresses (usually 8–16 bits in length). The IEEE 802.15.4 standard, for example, supports unique EUI-64 addresses carried in all radio chips, along with configurable 16-bit short

addresses. These networks by nature also support broadcast (address 0xFFFF in IEEE 802.15.4), but do not support native multicast.

IPv6 addresses are 128 bits in length, and (in the cases relevant here) consist of a 64-bit prefix part and a 64-bit *interface identifier* (IID) [RFC4291]. *Stateless address autoconfiguration* (SAA) [RFC4862] is used to form the IPv6 interface identifier from the link-layer address of the wireless interface as per [RFC4944]. For simplicity and compression, 6LoWPAN networks assume that the IID has a direct mapping to the link-layer address, therefore avoiding the need for address resolution. The IPv6 prefix is acquired through Neighbor Discovery Router Advertisement (RA) messages [ID-6lowpan-nd] as on a normal IPv6 link. The construction of IPv6 addresses in 6LoWPAN from known prefix information and known link-layer addresses, is what allows a high header compression ratio. 6LoWPAN addressing is discussed in detail in Chapter 2. A reference for IPv6, including the IPv6 addressing model, is provided in Appendix A.

1.3.4 Header format

The main functionality of 6LoWPAN is in its LoWPAN adaptation layer, which allows for the compression of IPv6 and following headers such as UDP along with fragmentation and mesh addressing features. 6LoWPAN headers are defined in [RFC4944] which has been later improved and extended by [ID-6lowpan-hc]. 6LoWPAN compression is stateless, and thus very simple and reliable. It relies on shared information known by all nodes from their participation in that LoWPAN, and the hierarchical IPv6 address space which allows IPv6 addresses to be elided completely most of the time.

The LoWPAN header consists of a dispatch value identifying the type of header, followed by an IPv6 header compression byte indicating which fields are compressed, and then any in-line IPv6 fields. If, for example, UDP or IPv6 extension headers follow IPv6, then these headers may also be compressed using what is called next-header compression [ID-6lowpan-hc]. An example of 6LoWPAN compression is given in Figure 1.10. In the upper packet a one-byte LoWPAN dispatch value is included to indicate full IPv6 over IEEE 802.15.4. Figure 1.11 gives an example of 6LoWPAN/UDP in its simplest form (equivalent to the lower packet in Figure 1.10), with a dispatch value and IPv6 header compression (LOWPAN_IPHC) as per [ID-6lowpan-hc] (2 bytes), all IPv6 fields compressed, then followed by a UDP next-header compression byte (LOWPAN_NHC) with compressed source and destination port fields and the UDP checksum (4 bytes). Therefore in the likely best case the 6LoWPAN/UDP header is just 6 bytes in length. By comparison a standard IPv6/UDP header is 48 bytes in length as shown in Figure 1.12. Considering that in the worst case IEEE 802.15.4 has only 72 bytes of payload available after link-layer headers, compression is important. The 6LoWPAN format and features are described in detail in Chapter 2. Note: these figures showing packet formats are in *box notation*, see Section A.1 for an explanation.

1.3.5 Bootstrapping

Applications of 6LoWPAN most often involve completely autonomous devices and networks, which must autoconfigure themselves without human intervention. Bootstrapping first needs to be performed by the link layer, in order to enable basic communication between nodes within radio range. Basic link layer configuration usually involves the channel setting, default

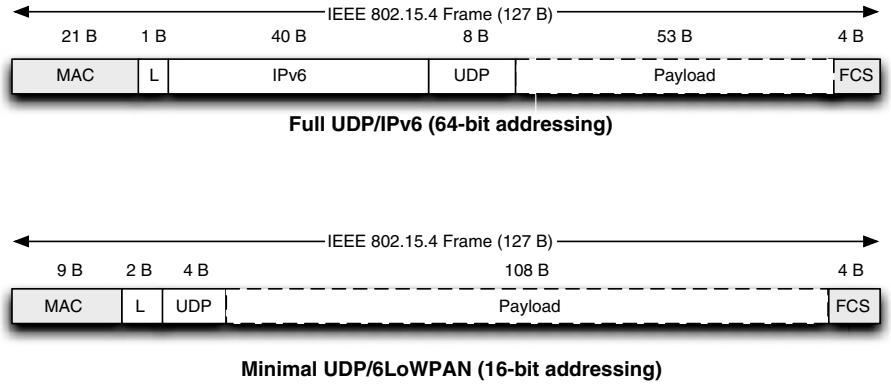

Figure 1.10 6LoWPAN header compression example (L = LoWPAN header).

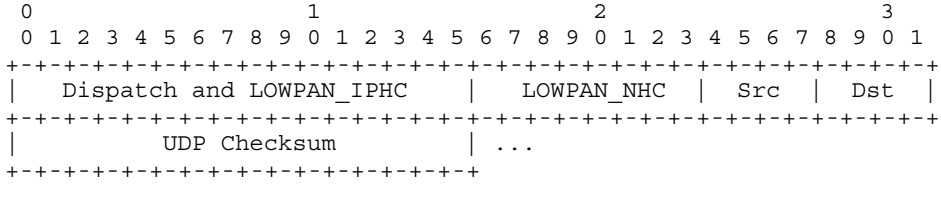

Figure 1.11 6LoWPAN/UDP compressed headers (6 bytes).

security key and address settings. Once the link layer is functioning and single-hop communications between devices is possible, 6LoWPAN Neighbor Discovery [ID-6lowpan-nd] is used to bootstrap the whole LoWPAN.

Neighbor Discovery is a key feature of IPv6, which handles most basic bootstrapping and maintenance issues between nodes on IPv6 links. Basic IPv6 Neighbor Discovery is specified in [RFC4861], but is not suitable for use with 6LoWPAN. The 6LoWPAN working group has defined *6LoWPAN Neighbor Discovery* (6LoWPAN-ND) optimized for low-power wireless networks and 6LoWPAN in particular [ID-6lowpan-nd]. The 6LoWPAN-ND specification describes network autoconfiguration and the operation of hosts, routers and edge routers in LoWPANs. A registry of the nodes in each LoWPAN is kept in the corresponding edge router, which simplifies IPv6 operation across the network and reduces the amount of multicast flooding. Additionally 6LoWPAN-ND enables LoWPANs covering many edge routers connected by a common backbone link (e.g. Ethernet), and the unique generation of short link-layer addresses. Chapter 3 looks at bootstrapping issues and Neighbor Discovery in detail. See Appendix A for a reference on basic Neighbor Discovery.

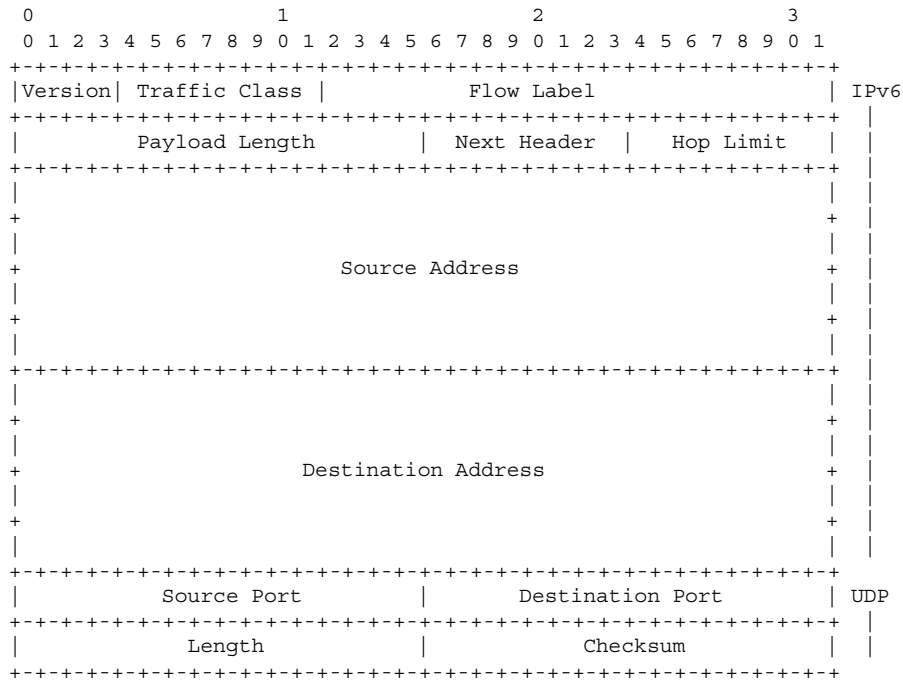

Figure 1.12 Standard IPv6/UDP headers (48 bytes).

1.3.6 Mesh topologies

Mesh topologies are common in applications of 6LoWPAN such as automatic meter reading and environmental monitoring. A mesh topology extends the coverage of the network, and reduces the cost of needed infrastructure. In order to achieve a mesh topology, multihop forwarding is required from one node to another. In 6LoWPAN this can be done in three different ways: link-layer mesh, LoWPAN mesh or IP routing. Link-layer mesh and LoWPAN mesh are referred to as *Mesh-Under* as the mesh forwarding is transparent to the Internet Protocol. IP routing is referred to as *Route-Over*.

Link-layer mesh is possible with some wireless technologies that include multihop forwarding features such as the recently completed IEEE 802.15.5 standard [IEEE802.15.5]. The original 6LoWPAN specification [RFC4944] includes an option for carrying mesh source and destination addresses, which can be used by a forwarding algorithm. No standard algorithms for use with this mesh header have been defined, and therefore the realization of LoWPAN mesh has been implementation specific. Currently, the most common technique instead employs IP routing. Routing with 6LoWPAN works just as with standard IP stacks, an algorithm updates a routing table which IP uses to make next-hop decisions. The Internet protocol is agnostic to the routing algorithm, and simply forwards packets. IP routing algorithms for mesh networking are developed in the IETF MANET working

group [MANET] for generic ad hoc networks, and in the IETF ROLL working group [ROLL] specific to wireless embedded applications such as industrial and building automation. IP routing issues and algorithms are discussed in detail in Section 4.2 including information on ROLL.

1.3.7 Internet integration

When connecting a LoWPAN to another IP network or to the Internet, there are several issues to be considered. 6LoWPAN enables IPv6 for simple embedded devices over low-power wireless networks by efficiently compressing headers and simplifying IPv6 requirements. Issues to be considered when integrating LoWPANs with other IP networks include:

Maximum transmission unit: In order to comply with the 1280 byte MTU size requirement of IPv6, 6LoWPAN performs fragmentation and reassembly. Applications designed for the Wireless Embedded Internet should however try to minimize packet sizes if possible. This is to avoid forcing a LoWPAN to fragment IPv6 packets, as this incurs a performance penalty. Additional considerations on fragmentation avoidance are covered in Section 2.7.2.

Application protocols: Application protocols on the Web today depend on payloads of HTML, XML or SOAP carried over HTTP and TCP. This results in payloads ranging in size from hundreds of bytes to several kilobytes. This is far too large for use with 6LoWPAN Nodes. End-to-end application protocols should make use of UDP and compact payload formats (preferably binary) wherever possible, as discussed further in Chapter 5. Technologies which are capable of the transparent compression of web services into a format suitable for 6LoWPAN Nodes are especially interesting.

Firewalls and NATs: In real network deployments firewalls and network address translators (NATs) are a reality. When connecting 6LoWPAN through these there may be several problems that need to be dealt with, for example the blocking of compressed UDP ports and non-standard application protocols used for 6LoWPAN applications, along with the unavailability of static IP addresses.

IPv4 interconnectivity: 6LoWPAN natively supports only IPv6, however often it will be necessary for 6LoWPAN Nodes to interact with IPv4 nodes or across IPv4 networks. There are several ways to deal with IPv4 interconnectivity, including IPv6-in-IPv4 tunneling and address translation. These mechanisms are typically collocated on LoWPAN Edge Routers, on a local gateway router, or on a node configured for that purpose on the Internet. IPv4 interconnectivity is covered in Section 4.3.

Security: When connecting embedded devices to the public Internet, security should always be a major concern as embedded devices are limited in resources and are autonomous. This is very much so with 6LoWPAN as node and network limitations prevent the use of the full IPsec suite, transport layer (“socket”) security or the use of sophisticated firewalls on each node. Although link-layer security inside a LoWPAN (employing the 128-bit AES encryption in IEEE 802.15.4) provides some protection, communication beyond LoWPAN Routers is still vulnerable. This increases the need for end-to-end security at the application layer. Security is dealt with further in Section 3.3.

1.4 Network Example

In this section we give a short example of how 6LoWPAN works in practice, concentrating on the basic things that occur during bootstrapping and operation. Figure 1.13 shows an example deployment of a Simple LoWPAN, connected through a backhaul link to the IPv6 Internet. The LoWPAN consists of an edge router, three LoWPAN Routers (R) and three LoWPAN hosts (H). Additionally there is a remote server on the Internet. This LoWPAN is based on IEEE 802.15.4 and uses IP routing (which is why there are LoWPAN Routers). Fake IPv6 subnet prefixes and addresses of nodes are included in the figure to make it easy to follow the example (in reality addresses would be longer).

The router to the Internet advertises the IPv6 prefix 2001:300a::/32 on the backhaul link, which is used by the edge router for autoconfiguration. The edge router then configures the IPv6 prefix 2001:300a:1::/48 to its IEEE802.15.4 wireless interface. Note that the LoWPAN and backhaul link are on different subnets as this uses the Simple LoWPAN model. The IEEE 802.15.4 wireless devices in the LoWPAN assume a default channel and security key settings. The edge router starts advertising the IPv6 prefix, which is used by the three routers to perform Stateless Address Autoconfiguration, and to register with the edge router

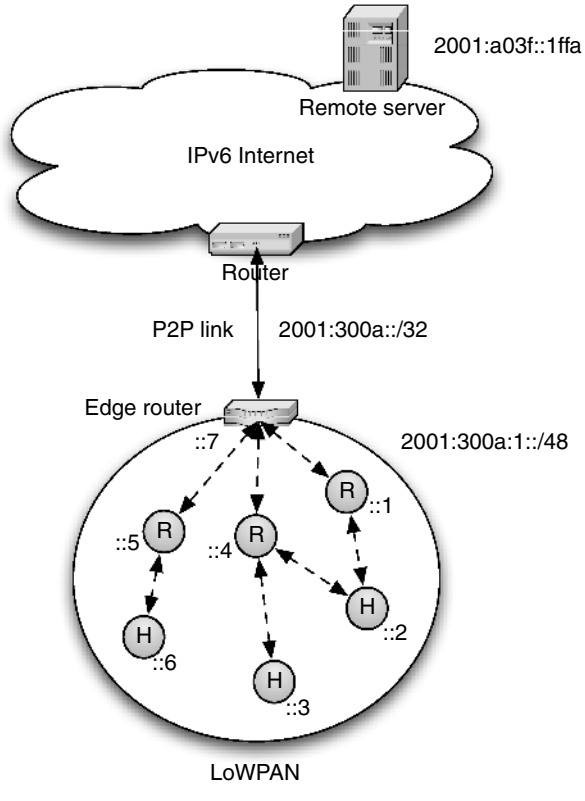

Figure 1.13 A 6LoWPAN example.

using 6LoWPAN-ND. Each LoWPAN Node now has an IPv6 address with a 64-bit IID and additionally receives a generated IPv6 address with a 16-bit IID from the edge router during registration. The IID part of the IPv6 address is shown in the figure as, for example, ::1, which has a full IPv6 address of 2001:300a:1::1. In reality these IIDs would be constructed from 16-bit random numbers. In turn the routers advertise the same prefix to the three hosts, which also register with the edge router. The topology inside the LoWPAN can change freely without affecting the IPv6 addresses of the nodes.

The Neighbor Discovery traffic used for advertising routers and registration is used to initialize the routing algorithm of the routers. Example routes are shown by dashed lines in Figure 1.13. IPv6 source and destination addresses of LoWPAN Nodes are elided during communication. A packet sent to a node on the same link (e.g. ::6 to ::5) does not require in-line IPv6 addresses at all as the link-layer header already contains the source and destination IEEE 802.15.4 addresses. If a packet is forwarded over multiple hops, then just the 16-bit source and destination addresses are carried in-line (e.g. ::3 to ::7) and those addresses are used for routing the packet. Packets destined outside the LoWPAN include either a full IPv6 destination address or a compressed one if compression context for that address is advertised in the LoWPAN. For example, LoWPAN host 2001:300a:1::6 may send a packet to the remote server at 2001:a03f::1ffa. The edge router expands the compressed LoWPAN and IPv6 headers to a full IPv6 header along with the UDP header if compressed. Incoming packets are also processed at the edge router, compressing IPv6 and UDP headers as much as possible.