Internet of Things: Concepts and System Design

Author: Milan Milenkovic, ISBN 978-3-030-41346-0

Chapter 1: Introduction and Overview

Internet of Things (IoT) systems connect the physical world to the Internet. Basically, IoT works by attaching real-world interfaces to the Internet, such as sensors that provide data and actuators that act upon their surroundings. In effect, IoT systems provide the technology and means to instrument, quantify and actuate the physical world.

Connecting sensors adds physical-world data, and in a sense awareness, to the Internet. This addition is a transformational change, it basically bridges the gap between physical and virtual/cyber worlds that has persisted since the invention of computing. In effect, IoT augments the Internet with all its features and capabilities by adding to it the real-world dimension. With the incorporation of IoT, Internet becomes a web of people, information, services and things, essentially the Internet of everything. As Kevin Ashton, who coined the term Internet of Things points out "In the twentieth century, computers were brains without senses - they only knew what we told them. In the twenty-first century, because of the Internet of Things, computers can sense things for themselves." [1]

There is no formal and commonly accepted definition of IoT. Many attempts tend to be descriptive and list system attributes, such as "a global infrastructure for the information society, enabling advanced services by interconnecting (physical and virtual) things based on existing and evolving interoperable information and communication technologies" [2] or "the extension of Internet connectivity into physical devices and everyday objects. Embedded with electronics, Internet connectivity, and other forms of hardware (such as sensors), these devices can communicate and interact with others over the Internet, and they can be remotely monitored and controlled. The IoT adds the ability for IoT devices to interoperate with the existing Internet infrastructure." [3]

Analysts predict that the installed base of connected devices will number in tens of billions in the next few years, and the number of sensors is projected to grow to hundreds of billions in the near future [4]. In the same time frame, the business impact of IoT is projected to be in trillions of US dollars, comparable to the GDP of the world's largest economies. Regardless of how projections turn out, those are staggering numbers indicating a truly transformational change.

Connection of real and cyber worlds using the Internet fabric and protocols, enables many new types of hybrid interactions and thus the potential for the creation of the plethora of exciting new uses, applications, and business models. Its development holds the promise to profoundly impact not only the industry but also many aspects of the daily lives and wellbeing of people.

IoT Systems

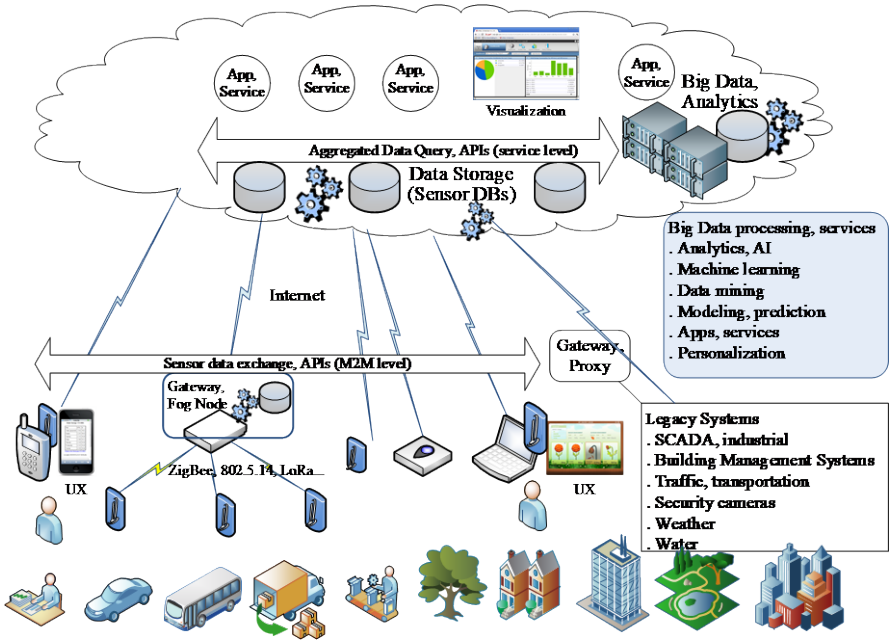

Fig. 1.1 IoT Systems

Figure 1.1 illustrates the major functional components of an IoT system. At the bottom are pictorial representations of some of the applications of IoT, including: office, automotive, transportation, manufacturing, agriculture, home, office building, and smart cities. The layer above it illustrates some representative categories of IoT sensors and things.

Some sensors and things are designed to connect directly to the Internet and communicate with applications and services residing in the cloud. Examples include a variety of devices, often marketed with a moniker smart, such as security cameras, fire sensors, thermostats, appliances and power meters.

Other sensors connect to the rest of the IoT system and the Internet using intermediaries, such as gateways. Gateways and fog nodes are usualy more powerful devices connected to collections of usually low-end sensors attached to it via local network links, often wireless, such as ZigBee, variants of 802.5.14 networks, Bluetooth and low-power Wi-Fi. Gateways usually provide wide-area connectivity and edge processing for the attached sensors that may come in the form of protocol conversion, data storage and filtering, event processing and analytics.

Billions of smart phones already in user's hands play multiple useful roles in IoT systems. They can act as smart connected sensors, gateways, and user interaction devices. Phones have a variety of built-in sensors - such as fingerprint, pressure,

light, Hall sensor, barometer, temperature, geomagnetic, accelerometer, gyro, proximity, and GPS global positioning system (GPS) - that can be used to report device location and state. These sensors are program-accessible by applications and therefore connected or connectable to the Internet. Many other location-based and augmented reality (AR) applications are possible by combining these with data from other phone sensors, such as cameras and microphones.

Smart phones and tablets may also act as a form of IoT gateways for other devices on their reachable local networks, such as multitudes of pairable Wi-Fi and Bluetooth enabled sensors including fitness trackers, smart watches, cars and home automation devices. In Figure 1.1, smart-phone types of devices are shown as being both connected sensors and user interaction devices. Smart phones can accept user input and visualize outputs from applications, provide user interfaces for data and control inputs, as well visualize any IoT data of interest available locally or anywhere else on the Internet.

IoT endpoints and sensors may also be embedded into the existing IT infrastructure and devices, such as personal computers. Adding IoT sensors and interfaces to existing IT devices, such as PCs and laptops, can provide coverage close to users and lower system cost by capitalizing on the host's power, connectivity, processing and storage resources. In addition, a PC can act as a gateway and a rich user-experience (UX) endpoint in an IoT system.

In addition to sensors and devices designed expressly for IoT uses, there are numerous sensors collecting real-world data in various forms of legacy systems such as industrial automation, energy and health systems, building-management systems (BMS), and supervisory and control systems (SCADA). For example, a BMS may contain thousands of physical-world sensors, such as temperature and motion in each room of an office building. Most of them are locked in proprietary formats, but BMS and SCADA systems are increasingly becoming interfaced to the Internet at their top levels of hierarchy for at least restricted forms of sharing and remote control. Data from such systems can be incorporated into IoT systems for aggregation and processing purposes with the aid of protocol and data format translators, often implemented as purpose-built gateways.

'Back-end" processing, often performed by remote servers in IoT systems, is depicted by a generic cloud in Fig 1.1. The picture is not intended to imply that there is a single centralized cloud in the IoT space, there are actually many, but rather to highlight some of the key components of an IoT system that support processing of data from large numbers of sensors by a variety of applications. Conceptually, the cloud is usually the top layer in an IoT system hierarchy, where data from a variety of diverse sources are aggregated and processed for optimization and discovery of global trends and relations. Depending on their nature and real-time requirements, incoming sensor data and events may be processed "in flight" as streams, stored for post-processing and archival purposes, or both. In addition to supporting and running applications, an IoT cloud may also contain some common services such as large-scale storage, analytics-processing engines, data visualization

and graphing, as well as management functions such as security and provisioning, not shown in Fig 1.

Machine learning (ML) and artificial intelligence (AI) algorithms are usually operated in the cloud where they can work with large aggregations of data. In current practice, such algorithms tend to perform better on large data sets that improve training and consequently their prediction capability. The value of IoT data aggregations increases significantly if they can acquire and store data in an interoperable format, thus supporting better analytics on larger data sets amassed from multiple sensing domains.

In terms of typical IoT data flows, sensors monitor physical-world status and provide the corresponding digital read-out of those, sampled at some data rate that may be fixed or variable through command settings. Flows of bits in IoT systems can be bi-directional, with sensor data typically flowing upstream to the cloud and towards applications, and control and status setting flowing downstream, such as a command to open a valve.

In addition to data capture, sensor streams need to be tagged with metadata to provide additional information about the meaning of data and the context in which they were acquired. Metadata may indicate the nature of the data, e.g. temperature, reporting location, ownership and structural relations. Metadata may be added to sensor data streams anywhere on the path from the point of capture to sensor data storage in the cloud and in different stages of the lifecycle, including installation, provisioning, operation and management.

Although not explicitly indicated in Figure 1.1, applications that process and react to sensor data may reside at the edge, in the cloud, or distribute their functionality across multiple system components. For instance, implementation of a fitness application can be a fairly complex IoT system. It starts with a wearable device that may include a variety of sensors to detect a user's activity, motion, position and perhaps some vital indicators such as heart rate or EKG. This can be a complex embedded system that needs to constantly acquire and pre-process data, have a long battery life and be as small as possible. It usually contains some preliminary processing functionality, such as data capture, pre-processing, time stamping, and buffering of data until its transmission to the phone or cloud is completed. Additional computation and storage may be placed in the portion of the application running on the smart phone that the sensor is paired with. In addition, the phone application can provide a user interface to visualize the data and customize settings such as individual targets. The phone also typically acts as a gateway by relaying user data to the cloud portion of the fitness application and service that can archive user data and provide more comprehensive types of analyses, such as trending and comparisons of results with a user's friends groups. In addition, the cloud application can aggregate hopefully anonymized data from its entire user base that may be used to study population trends for various uses, such as user reference and medical studies.

Figure 1.1. illustrates data and control flows related to the primary functional aspects of operation of IoT, namely collection of data and processing to act on it.

This is sometimes referred to as the data plane or user plane. Another important dimension of an IoT system, not shown in Figure 1.1., is the control plane that maintains security and operational aspects of the IoT infrastructure itself. Its primary components are security and device management that include addition and bringing up of new devices and operating the system safely and reliably through constant monitoring and remediation when necessary.

Why Now?

A confluence of technological and infrastructure developments centered around the Internet form much of the basis and impetus for the construction of IoT systems. The major ones include:

- Industry 4.0 and digitalization
- Sensors – installed base, variety, lower cost, easier integration
- Smart phones – sensors, gateways, UI devices
- Cloud computing – global, capacity on demand
- AI and ML technologies – actionable insights with IoT data
- Internet – technology, global infrastructure, users

Industry 4.0 generally refers to the digital transformation of manufacturing using a combination of technologies including automation, robotics, digital-to-physical transfers such as 3D printing, big data, ML and AI, robotics, and augmented reality (AR). Some of its key foundations are machine-to-machine communications, Internet compatibility and interoperability. With sensing and digitization of data as its starting point, Industry 4.0 shares many ingredients with Internet of Things, and IoT is generally regarded as one of its foundational technologies. Enterprise-scale investments in Industry 4.0 will facilitate deployment and accelerate the industrial dimension of the IoT technology evolution. One of the aspects of Industry 4.0 is bridging and format translation to incorporate the data already being collected by the large installed base of sensors in legacy systems mentioned in the previous section.

The availability of a variety of connectable *sensors* in the consumer space facilitates construction of IoT systems and enables many potential applications. Home automation and monitoring devices, such as smart thermostats, lights, security systems and Internet-enabled cameras, can perform useful functions in isolation and be combined to perform more complex functions by engaging in coordinated ensemble behaviors. A growing base of personal sensors, such as wearables, enables the continuous tracking of physical activities and vital signs for fitness and health applications.

Another major contributor of inexpensive and widely available sensors that are comparatively quick and easy to deploy is the rapidly growing “maker” movement. [5, 6]. It has brought about inexpensive processor boards for prototyping IoT

systems, such as Raspberry Pi [7], and commonly used and de facto standardized I/O interfaces, such as GPIO (General Purpose Input Output) and packaging in form factors such Arduino shields [8]. This has led to the availability of a wide range of compatible sensors and actuators – such as temperature, light, humidity, pressure, presence, air quality, touch, distance, current, vibration, dust, heart rate, sound, switches, relays, motor drivers – whose cost tends to be orders of magnitude lower than their traditional industrial counterparts [9]. These sensors usually come with drivers and libraries for popular development environments, such as Arduino. Prior to these developments, constructing of a hardware prototype to test an idea with potential users required costly investment into custom development, capital expenditures and delays in scheduling the generally less profitable small production runs of product prototypes. The availability of low-cost sensors, platforms and software tools drastically reduces the time, costs and risks of development, testing, and experimentation with IoT prototypes. This spurs innovation by allowing rapid and inexpensive cycling and testing of variants to validate a product and its features before committing resources to its production.

As described in the previous section, *smart phones* provide a large installed base of smart and connected sensors that can be used in a variety of applications. Phones can also act as gateways between sensors in the user's proximity and applications and services residing locally or anywhere in the global cloud. In addition, phones provide user interfaces for data and control inputs, as well as visualization of virtually any IoT data and relevant content on the Internet.

The availability of *cloud computing* facilitates the construction and deployment of an IoT system with global reach and without the requirement for up-front capital investment in servers, cloud networking and storage. Moreover, it comes with the appearance of almost instant and limitless scaling of capacity to meet variations in demand. Cloud technology and commercial offerings provide the capability to meet IoT system requirements in terms of volume and bandwidth for large-scale data ingestion, real-time event processing and Internet-scale databases and storage for large aggregations of data.

Advances in *AI and ML technologies* have led to major successes in applications such as image processing and natural language understanding and translation. There are growing expectations that similar achievements may be accomplished by applying AI and ML techniques in IoT systems to obtain insights leading to actionable outcomes. There are many AI and ML tools, some of them with libraries of neural network configurations with pre-trained parameters for specific applications that may be used as a starting point to accelerate IoT experimentation and adaptation. Many of these tools are available on the Internet and hosted by cloud providers.

The most important component of IoT systems is their foundation and key pillar, the *Internet* itself. We use the term Internet in a broad sense that includes not only its original underpinning of IP-enabled networks and protocols, but also the world-wide web and numerous applications that it has enabled. Internet size, vibrancy, global reach, services such as public clouds, users, development tools and an army

of developers provide a formidable foundation for IoT systems to add to and build upon.

The Internet provides a set of standardized communication protocols and data markup languages that allow the exchange of data by endpoints and their use by multiple applications. Standardization and adherence to common specifications enable independent design of components that are modular and interoperate. As a result, a user can select a device of their choice from one vendor, pick a browser from another source and use the combination to access Internet-compliant content posted anywhere in the world. The content itself may be created using any compliant authoring tool supplied by a variety of vendors and hosted on any of a variety of web servers that may run on any of a number of different hardware and software platforms.

Internet technology has been proven in the construction and operation of immensely scalable, (almost) always available systems with the ability to handle millions of servers and billions of users while providing global reachability to both producers and consumers of data. Moreover, the huge size of its already installed base and global infrastructure that is already deployed provides a nearly ubiquitous availability of Internet in many of the populated parts of the world. Its modularity and adherence to standards has spurred innovation and created a competitive marketplace with many hardware and software offerings that can be mixed and matched to create services and solutions.

How IoT Systems are Different?

While being able to capitalize on much of the technology and infrastructure of the Internet, IoT systems have some characteristics that are unique and structurally different and require specialized design and engineering to implement them. Much of the rest of this book is devoted to those topics.

Some of the key IoT differences and challenges include:

- Information models, standards, data and metadata formats and semantics
- Search
- Topology, inverse client to server ratios, and reverse data flows
- Continuous, time-series data streams
- Distributed function placement – edge, fog, cloud

In order to illustrate some of the differences, consider a simplification of the common usage mode of the Internet to access a web page from one of the billions available on the web. The (human) user may start by searching for a phrase of interest or a partial name of the source. The search engine of choice typically returns a number of clickable links (URLs) to pages that contain content that is likely relevant. The page chosen by the user is displayed with the appearance and format

created by its authors. It can also contain active links to other relevant pages that the user can read and navigate.

An important objective for IoT systems is to enable something of comparative functionality and power for IoT data. In order to provide a similar degree of flexibility that we came to expect on the Internet, IoT systems need to grapple with the difficult problem of machine-level understanding, i.e. interpretation of data generated by sensors and things.

Note that in the previous description of the Internet, data are destined for consumption and interpretation by humans. The Internet provides a method and protocols to transport data from servers to clients and use some annotation, such as HTML, to annotate the data for proper rendering at the receiving end. Rendering implies displaying the text with the author-specified appearance, such as the font type, and proper placement of images and other types of media. But beyond essentially mechanical rendering, there is no semantic interpretation of the data by the receiving client. The meaning of the data is defined by the natural language in which the text is authored. If you do not understand that language, no semantics is conveyed. As an example, look up a web site of a major newspaper in a language that you cannot read or do not speak (in my case, Japanese or Arabic). More precisely, the textual Internet provides the syntactic and the structural interoperability between the endpoints, but not the semantic one.

In IoT systems, data exchanged between machines at endpoints are sensor readings and actuator commands. These are typically device- and domain-specific numbers such as temperature readings, with annotations such as the nature of measurements and units of measure, and control strings that do things like change a set point. Applications (not humans) at the receiving end need to be able to "understand" and interpret messages in the payload in order to process them appropriately. For this to work, there has to be some common understanding between the senders and the receivers on what the data means, and how to encode and annotate it for transmission and to correspondingly decode it upon reception.

The problem here is that there is *no commonly accepted convention for representing data in IoT systems*. In its absence, it is not possible to independently develop IoT endpoints that would interoperate as is common on the rest of the Internet. A number of standardization efforts are under way to address this problem. The challenge faced by IoT standards is to replicate the Internet benefits of modularity and interoperability to data and control exchanges at the edge. This means that components may be developed independently, on different platforms and by different establishments, and reasonably expect to be able to connect and interoperate with others provided they correctly implement the applicable standards. IoT data representation is an important design problem and we devote two separate chapters of this book to IoT data definition, interoperability, and work in standards that are addressing it.

Internet *searching* techniques are not directly applicable to IoT since its "content" is not the textual information in a natural language that can be indexed by the common web crawlers. Consequently, it is not possible to search the IoT space for

content based on names or attributes of the endpoints. Moreover, many of IoT endpoints do not have Internet resolvable names for reasons of scale, domain-name assignments and, in some cases, are kept intentionally private for security. Such systems may deploy mechanisms of device self-description and discoverability or maintain dedicated device directories to enable limited, domain-specific searches. In any case, in order to enable attribute-based searching, IoT data need to be annotated by contextual and semantic information, usually in the form of metadata.

Traditional Internet servers tend to be located closer to the core of the network for higher bandwidth and to reduce access latencies. Clients with browsers tend to be more towards the edge of the network. The current statistics suggest that the number of servers on the Internet is on the order of 100 million, and the number of clients is closer to 4 billion. The bulk of data flows tends to be from the cloud core towards the edge as servers deliver requested data and media streams. This is partially evidenced by the fact that many edge connections, such as residential Internet links, tend to have asymmetric speeds, with downstream bandwidth often exceeding the upstream by an order of magnitude.

IoT systems are also different in terms of *topology, client to server rations, and primary direction of data flows*.

IoT nodes that source IoT data are referred to as servers. In practice these are usually the edge nodes and things that may have constrained capacity, power, connectivity and bandwidth. Clients are applications and services that consume IoT data. They may reside on peer edge nodes, in gateways, in fog nodes, or in the cloud. Thus, unlike on the Internet, clients in IoT systems may have higher processing and storage capacity then IoT servers.

In an IoT system, the total number of servers with data sources, i.e. billions of endpoints, largely outnumber the clients, i.e. IoT services and applications that consume them. An IoT server will typically provide data to a few clients or publish them to a messaging broker where any number of subscribed clients can use them. IoT clients implemented as cloud services or data aggregators may consume data from numerous IoT servers. The bulk of IoT data flows is from the edge servers towards the processing clients that are located closer to the system core.

Thus, in comparison with the mainstream Internet, IoT systems tend to have inverted client to server ratios and reverse data flows, from the edge towards the more centralized parts of the system where applications and services (clients) that operate on combinations and aggregations of data tend to reside.

Much of IoT data have *time-series* characteristics in the sense of being consecutive, time-stamped samples of sensor readings. Applications tend to retrieve such data by combinations of stream identities and periods of time, such as within the last hour or at a specified date. Volumes of IoT data can be very high as they are generated by machines and often continuously sampled, sometimes with very high frequency. IoT data may also have real-time processing requirements that may impose tighter bounds on latency tolerances. Thus, IoT system architecture and implementations need to account for potentially high volume and throughput of data

ingress from the edge. In addition, they need to meet processing and storage requirements of time series and real-time data.

Placement of functions is the quintessential design decision in distributed systems, such as IoT, where data and processing in the form of services and applications may not be co-located. Basically, it is a tradeoff that revolves around bringing the data to processing or processing to data. Considerations driving the decision include processing capacity, storage, latency, connectivity, bandwidth availability and cost. In IoT systems additional considerations may include real-time and local autonomy requirements in the sense that some degree of control functions should remain operational even when cloud access is not available.

In IoT applications where data volumes are high but significance of much of it is low, it is often better to process the raw data close to the source. For example, recognizing a human shape in a surveillance application may require constant acquisition of a video stream with potentially high bandwidth. However, since only fragments of interest are the ones with human shapes in them, considerable reductions in bandwidth may be realized by processing the video stream at the edge and forwarding data to the cloud only when the local AI infers a significant event, i.e. detection of a shape. In this example, potential savings in bandwidth have to be weighed against the cost and complexity of placing and managing processing capacity at the edge.

In principle, in IoT systems processing and storage functions may be placed anywhere along the data path, from the acquisition at the edge and traversal of intermediate processing and staging nodes, such gateways and fog, all the way to the cloud. Moreover, processing modules may be distributed at various points on the data path to form the processing pipelines to complete increasingly more complex functions. Judicious placement of processing and storage functions is one of the important factors in IoT systems design.

Value and Uses of IoT

What are the benefits and value that can be expected to be realized by deploying and using IoT systems? A somewhat oversimplified answer is provided by the old management adage – what you can measure, you can improve. IoT enables continuous sensing and measurement of the state and behavior of the physical world in order to react and ultimately improve on some of its aspects that can be controlled. Pervasive and continuous monitoring provides insights and a quantified view of the physical world. The extent and type of the resulting influence varies considerably depending on the nature of the system under observation. A manufacturing operation may be directly controlled by an IoT system using guidance provided by an AI system. On the other hand, continuous monitoring of a person's EKG via a smart personal sensor can identify and record anomalies and changes in vital signs that

preceded it to be made available to a physician for diagnosis informed in a manner that was previously not possible.

Two primary areas of potential application of IoT systems are in business and consumer segments. Some of the early IoT consumer applications are already becoming familiar – such as self-driving cars, wearable health and fitness sensors, and Internet-enabled home security cameras and smart thermostats. However, the major financial impact is initially expected in business uses. In business processes, use of IoT can lead to transformational changes that impact efficiency and profitability of enterprises, such as improvements in predictive maintenance, better asset utilization, and higher productivity.

In the remainder of this section we describe several applications to illustrate some of IoT uses that are under way or being contemplated. Many other current and potential uses of IoT systems are described, some in considerable detail, in other chapters in this book.

According to analysts, the ranking of IoT uses based on projected business value and financial impact may look as follows:

- Production environments - factories
- Cities
- Human health and fitness
- Retail environments
- Transportation and automotive
- Oil and gas, mining
- Utilities and energy
- Home
- Offices

Of course, projections and actuals may vary based on how various uses are classified and clustered, but our interest here is in identifying the major broad categories.

Production environments with repetitive work routines are at the top of the list because they can use digitalization to quantify their performance in terms of measurable outputs. Many of them already use automation and control systems to a certain extent. Adding IoT in such environments can improve the relevant key performance indicators and demonstrably justify the return on investment. Such facilities may have thousands of sensors that operate in smaller control loops to manage specific machines or production lines, using different and often proprietary data formats and protocols. IoT systems may be used to break those control and management silos and provide system-wide insights that can lead to global optimization and improved efficiency of a production site or ultimately the enterprise.

In addition to quantified insights into the facility, near-term uses and benefits of IoT in production environments include anomaly detection, preventive maintenance, and fault prediction based on combinations and correlations of sensor data. The Internet connectivity allows a manufacturer to aggregate data from large segments or their entire installed base of products to create data-based profiles and to

provide expert operational guidance to its customers and their IoT systems. It also makes it possible to track the rate and nature of defects for design improvement purposes.

Cities can benefit from the use of IoT to improve their operation and services in diverse areas, such as transportation and parking, street lighting, energy consumption, sustainability, and public health and safety. The grand vision is to provide holistic insights and coordinated actions across systems in a city. In practice, this will require instrumenting and automating of individual systems, and achieving interoperability to share collected data for global insights and control. There are many projects and much experimentation with creation of smart cities that should identify the most promising uses and yield tangible benefits when fully functional systems are deployed.

Human health and fitness applications often work in conjunction with wearable devices that are capable of constant monitoring of user activities and some vital signs. This provides a range of potential uses, from tracking and encouraging exercising to detection and recording of irregularities that can led to data-informed diagnostics and personalized treatment. The potential for continuous monitoring and detection of anomalies at the moment they occur, is a significant qualitative departure from the current practice of measuring of vital signs infrequently or when visiting the physician's office. Patients with chronic conditions can additionally benefit from continuous monitoring of the effectiveness of their treatment and intake of medications, as well as inform the physicians about patient's adherence to the prescribed regiment. Long-term potential impact of IoT on the practice of medicine can be truly transformational as "quantified self" provides detailed and continuous data on an individual. This in turn allows personalization of treatment based on rich and relevant data, as opposed to being guided by broad averages based on sparse and randomly sampled data in medical offices and studies, as is the norm today.

In *automotive uses*, self-driving and autonomous vehicles are probably the most publicly visible instantiations of IoT systems. They are based on extensive visual and ranging sensing of a vehicle's surroundings processed by AI-driven algorithms that produce controls and commands for autonomous navigation. Due to latency and autonomy requirements (the car should not crash if the cloud becomes unreachable), much of this processing is done locally at the edge, within the vehicle itself. However, data from perceived situations and outcomes, both favorable and unfavorable, are sent to the cloud and aggregated to refine the navigation algorithms and update the local versions executing in vehicles. Moreover, data gathered from individual vehicles may be used to continuously refine and update mapping information that is useful to drivers and essential to autonomous navigation. A car manufacturer can use data from its connected cars to get insights into typical usage patterns, type of travel and distances driven, battery usage for electric vehicles, types and rates of anomalies and failures and the like. All of this information provides valuable guidance for product improvements and design of future models, as well as providing a basis for a variety of additional services, customer loyalty programs, and targeted offerings. One new service being considered for autonomous vehicles is to make

them available for rental or ride sharing when not in use by their owners – by driving themselves to and from their service destinations.

Research and development of vehicle-to-vehicle networks is under way with the expectation that it will improve safety and efficiency by alerting nearby drivers to road hazards, such as treacherous road conditions, e.g. ice and fog, traffic accidents and congestions. This can improve public safety and traffic management through aggregation of data about and from the vehicles moving in and out of areas of interest. Other obvious and more incremental applications of IoT technologies in automotive systems include fault prediction and detection for maintenance and operational purposes.

In retail applications, the focus tends to be on user experience, real-time tracking of user activities and inventory for data-driven management of the supply chain.

Oil, gas and mining are usually characterized by equipment monitoring in geographically dispersed areas, with harsh conditions and lack of access to wired infrastructure, such as electrical power and networking [10].

Similar geographically dispersed and poor infrastructure conditions exist in agriculture uses when used for monitoring of soil condition, crops and livestock. Obviously, the applications and services are very different.

Transportation and logistics systems include a variety of applications centered around tracking goods in motion and fleets of ships and trucks that deliver them.

Utilities are expected to benefit from IoT in production and distribution via a smart grid, and consumption and load monitoring via smart meters. Other major potential improvements include the ability to measure the load in real time and adjust production accordingly. Even more interesting is the possibility to sculpt the load, i.e. activate or delay schedulable consumers – such as EV chargers and appliances – to coincide with the production cycles of renewable but variable energy sources, such as wind and solar. Adaptive matching of supply and demand by means of load shaping is useful because storage of energy is not yet commercially cost effective and mismatches result in waste or outages.

In *homes* the current focus tends to be on automation of chores and maintaining user convenience and comfort through home automation for security and ambient control such as lighting and temperature, smart thermostats and smart appliances. Voice assistants that can access the Internet and control some home devices are becoming popular in some markets. They can act as input, output, and control point or gateway for other home devices. Unfortunately, the use of different standards leads to fragmentation and general inability to control home devices from different manufacturers as a collection with coordinated behaviors to improve user experience.

In general, changes resulting from introducing IoT capabilities can be (1) transformational changes that improve the effectiveness of a system, and (2) new uses and applications that may enable new business models. IoT uses and applications described earlier combine components of the two to various degrees.

One of the features of IoT systems that enables some interesting new service and business models is the ability to continuously and remotely monitor the physical

world. Coupled with the Internet connectivity, it is giving rise to what is often referred to as “something as a service” model. One of its incarnations consists of providing the physical equipment with monitoring and control on an ongoing service basis, instead of an outright purchase. For example, in an engine or power as a service scenario, aircraft engine manufacturers may provide engines to airlines and offer continuous monitoring of their operational performance in exchange for service fees that may include components based on usage, such as operating hours or segments flown. For the manufacturer, benefits can include the constant recurring service revenue flow and the real-time insights into their products in operation across the customer base. Monitoring data can be used for early detection of anomalies and to improve maintenance. Aggregations of data can be used to improve operational analytics and AI algorithms. In addition, they can provide insights and learnings to guide product improvements and future evolution. For airlines, benefits can include expert monitoring and maintenance guidance, as well as conversion of the up-front capital expenses for engine acquisition into operational expenses that can be covered from the ongoing revenue stream. This can reduce the need for capital loans and looks much better on the corporate performance sheet.

The presented examples illustrate that the possibilities brought by the IoT systems and by introducing the real-world dimension and awareness to the Internet have significant commercial and transformational potential. However, we are very early in the development and deployment cycles of IoT technology and – if history of computer applications is any guide - its true potential remains to be realized by the future applications yet to be invented.

Issues and Challenges

IoT system implementations and their applications face a number of challenges that need to be addressed in order for the full potential of IoT can be realized. The major ones include:

- Technology and standards
- Security and privacy
- Business value and adoption

Technological challenges include a system design that is suited to the operational characteristics of IoT systems and addresses its differences from the traditional Internet. Predominant among these is the need to achieve interoperability of IoT data across specifications and domains. The current fragmentation of standards and proprietary approaches create silos that limit usefulness and the ability to harness the benefits of big and diverse aggregations of IoT data. One major analyst study [4] estimates that achieving interoperability can increase the potential IoT business impact by 40%. In its absence, a great opportunity remains in danger of not being realized.

Security is important in all computer uses, and especially in enterprise and production environments where breaches can result in significant financial and reputational losses. It is even more important in IoT systems that can directly impact the physical world where cyber-attacks can cause not only major damage and disruption, but even endanger the safety of people and the environment. The problem tends to be further aggravated in IoT systems that are connected to legacy operational systems, such as utilities and manufacturing, that may have inadequate and ossified security protections. In the consumer space, individual devices and home installations are often insecure and may create security and privacy exposures due to device design and configuration deficiencies resulting in inadequate security or one that deteriorates over time due to lack of updates and ongoing security management.

Privacy and the ability to safeguard user data are also of major importance in IoT systems, since their data streams can provide considerable details and real-time information on user activities and location. Unchecked revelation of such data may be exploited for not only unwelcome and potentially intrusive surveillance but also for nefarious purposes such as break ins or theft when users are known to be away.

Finally, large-scale commercial success and wide deployment of IoT in industry depends on creating uses, applications and services that create value. While experimentation with IoT and proofs of concepts abound in industry, one of the key challenges is to demonstrate tangible business value. This means identifying applications with considerable revenue potential that produce positive business outcomes and measurable improvements, such as increased efficiency, user experience, reduced costs, or some combinations of those.

IoT Systems: A Reconnaissance Flyover

This section provides a brief overview of the entire IoT system in order to identify key functions, components, and how they relate to each other. The presentation is provided from the two different points of view to highlight the purpose and structure of various parts of IoT systems (1) functional view, and (2) infrastructure view. The functional view focuses on what needs to be done, i.e. key processing stages of IoT data on their way from capture, to processing and actuation in some form. This is followed by the infrastructure view that focuses on where and how things get done, i.e. the key infrastructure and hardware components that execute those functions.

IoT System Functional View

The primary purpose of an IoT system is to collect real-world data and make them available to services and applications in order to create insights and act upon

them by affecting the real world in some way [11]. Implementation of those functions requires an infrastructure to run them and control functions to keep them secure and operational. In this section, we focus on the functional aspects and cover other views in the subsequent ones.

Figure 1.2 depicts a highly abstracted functional view of an IoT system with focus on data flows and types of processing from capture to output actions. It highlights the key stages in IoT data and control flows (1) data collection, (2) processing, and (3) acting upon the world based on the outcomes.

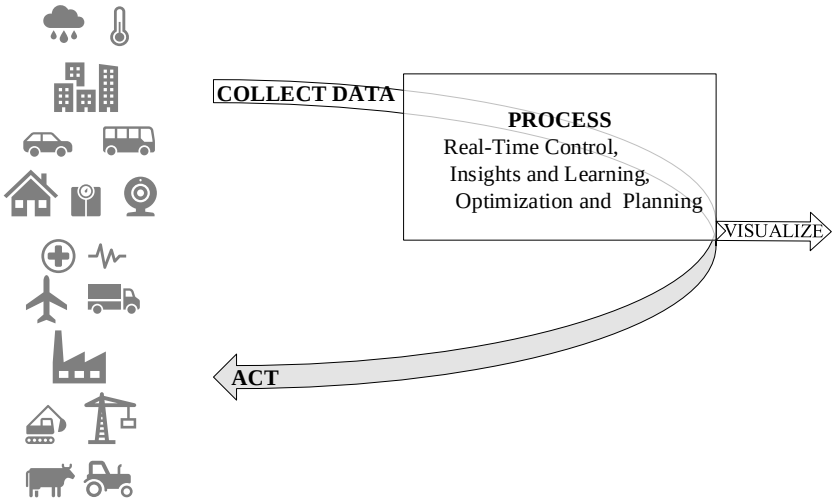

Fig. 1.2 IoT System Functional View

Data Collection

Data collection starts at the edge, with a sensor acting as a physical-cyber interface that monitors and reports states of some physical entity or device. The intent is to produce a digital representation suitable for use in the cyber space. This process may involve many practical details, such as signal conditioning, analog-to-digital conversion, scaling and conversion to engineering units for subsequent processing. These operations are important for the implementation of IoT systems and we cover many of them in subsequent sections and chapters. However, from the functional point of view, data collection results in digitized state samples of the physical world suitable for processing by the applications and services in the cyber domain.

Relatively early in the process, metadata needs to be captured and used to annotate the data. In IoT systems, metadata generally describe the nature and context of data capture, such as the sensor type, its location, and in some cases structural relationships to other elements of the system.

Data Processing and Visualization

Types of IoT data processing range from the simple control-loop algorithms performed on the incoming streaming data as they arrive, to the sophisticated forms of analytics and machine-learning algorithms that operate on combinations of streaming and archived data, events, and records of past behaviors and observations of the system.

Various steps of data processing may be implemented in increments or in entirety in the different components of IoT systems. In general, their scope and complexity tend to increase in the higher levels of system hierarchy, where more processing, power, storage and larger aggregations of data are available.

Common data pre-processing steps may include filtering, aggregation and comparisons to detect if the sampled data are in a special condition that warrants additional action, such as creation of an event or notification. When detected, they are forwarded to the processing services that may include some combinations of event-driven control loop algorithms, user notifications, or operator alarms.

The next level of sophistication in data processing is to provide optimizations and predictions of system behavior based on its current state, past behaviors, and guidance form algorithms, such as analytics and machine-learning models. Subsequent stages consist of learning system behaviors in order to transform that knowledge into effective insights and control actions.

In industrial and complex control systems, it is customary to visualize the system state and points of interest to system operators. These visualizations normally include data on the vital indicators of system state and notifications and alarms when faults or anomalous behaviors are detected. Operators usually have the option to zoom in and inspect any data point of interest for information and analysis purposes.

Improvements due to IoT deployments usually happen in stages. The main phases typically progress through the following stages (1) data collection and visualization, (2) insights and learning, and (3) optimizations and actions.

The first stage is to instrument the system under observation by installing and connecting sensors and devices with the rest of the IoT system. The resulting monitoring of collected data leads to insights and evidence-based understanding of the observed systems. While algorithmic analysis is often the goal, a somewhat unappreciated but important aspect of IoT instrumentation and continuous data collection is that they can immediately provide valuable insights to system operators. Human operators familiar with the system can use their experience and natural intelligence and act as very good “inference engines” in generating insights that they can deploy for more effective management of system operations. Moreover, they can also identify areas of potential improvements for analytics to focus on.

The third stage is to deploy the predictive and prescriptive forms of analytics and AI algorithms, informed by the prior human and machine insights and learnings.

Acting

Acting upon insights and predictions is the output and the ultimate purpose of deploying IoT systems. Actions can take different forms, from simple remote actuation initiated by operators in response to visualized conditions in a basic monitoring configuration, to automated guidance of control points that proactively manages conditions in a smart building to maximize user comfort and optimize energy efficiency. Actions can be implemented as direct actuation or indirectly, in the form of advice to system operators or optimizations resulting in adjustments to the manufacturing process. They can also include identification of cause of failures and anomalous conditions followed by direct or indirect execution of the appropriate remediation actions.

With the addition of interoperable data formats and Internet connectivity, the scope of data aggregation in IoT systems can grow to potentially any level including multi-domain systems such as smart cities and regions. Such data may be used for detecting and acting upon regional and even global insights. For example, one could aggregate behavioral and energy efficiency data from all IoT managed buildings in a large region and use it to train and improve the AI and ML algorithms for cross-domain effectiveness in all of them.

IoT System: Infrastructure View

This section covers the IoT system infrastructure components, including data processing, storage, and communication, that host and execute its functions outlined in the previous section.

Large IoT installations can be complex distributed systems with a variety of components and multiple levels of hierarchy. Figure 1.3 illustrates some key infrastructure components of an IoT system. Edge components are depicted towards the bottom, the communications layer mostly in the middle, and upper levels of system hierarchy ending with the cloud are shown on top.

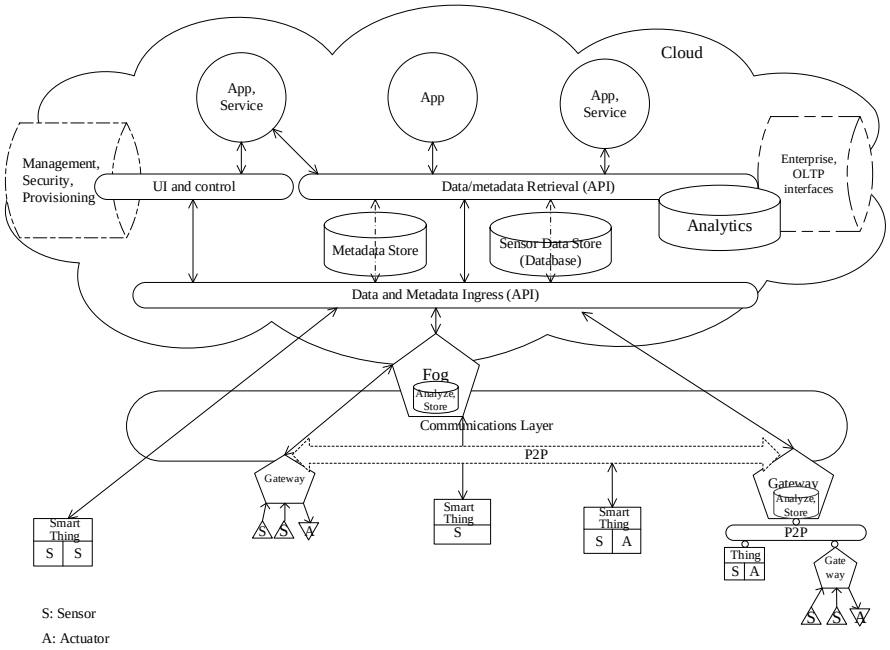

Fig. 1.3 IoT System Infrastructure View

Edge components include sensors, smart things, gateways and fog nodes. The communication layer provides connectivity among system components that they can use for horizontal peer-to-peer interactions within a level of system hierarchy or for cross-level communications towards the Internet and the cloud. The cloud portion depicts the back-end part of the IoT infrastructure where large-scale data aggregation and processing take place.

Edge Components

At the edge layer, Figure 3.1 depicts several types of components, including smart things and gateways. Smart things are generally characterized as devices that contain sensors or actuators and contain sufficient functionality to perform some local operations, connect to the Internet to report their data or receive actuation requests, and interact with services. Smart things in Figure 1.3 are shown as connecting directly to the cloud, to a fog node as a processing intermediary, and to other network nodes, such as gateways, in peer-to-peer configurations.

In practice, communication is carried at the behest of software agents that implement the relevant functions and services at the sending and the receiving party. If both parties reside on the same network or are at the same level in system hierarchy, this communication is referred to as peer-to-peer (P2P), and sometimes more

generally as machine-to-machine (M2M) communications. Technically, all communications between IoT devices themselves as well as with the applications and services are of the M2M type, so we use the term P2P to depict peer level communications, which are shown as horizontal paths in Fig 1.3, as opposed to the edge to cloud which would follow the vertical paths. Depending on the underlying physical network, a single node may be able to engage in both types of connections.

All of these directions are conceptual in the sense of indicating the hierarchical position of endpoints in a given communication. With the Internet as the underlying network layer, any addressable point can connect to any other addressable point anywhere in the world, subject only to authorization.

P2P communications can be somewhat distinct in the sense that they can take place between nodes using simplified or application-specific protocols that are not Internet compatible, such as Bluetooth or legacy industrial devices. This can reduce the load and simplify design of the constrained nodes. This is illustrated in the lower right corner of Figure 1.3 that depicts a cluster of devices, a thing and a gateway connected to each other and another gateway in a P2P manner. As long as they support a common set of private protocols, the three nodes can engage in P2P communications. However, to qualify as IoT endpoints, at some point in the system, say a gateway, a transition to the Internet compatible protocols needs to be made in both directions, so that functional exchanges of messages can take place with the authorized endpoints anywhere on the Internet.

Gateways are edge devices in IoT systems to which one or more basic sensors are connected and dependent upon for wide-area connectivity and optional additional services. Their basic role, as the name implies, is to provide connectivity between locally connected sensors and the Internet.

Gateways can reside in physical structures, such as a building, where an Internet access point already exists and the gateway needs to be connected, for example via an Ethernet connection. If a gateway is not near a usable existing Internet access point, depending on the distance to such a point, it may provide a medium-range link, such as LoRa, or a wide-area link to the Internet, say via a private wide-area network or a telephone company link such as GPRS or LTE.

The early implementation of IoT systems tended to place most of the processing, analytics and storage in the cloud. In such settings sensor data are sent to the cloud directly or via predominantly communications gateways with minimal intermediate processing. As discussed in detail in Chapter 2, this approach has drawbacks that include longer communication and processing latencies, potentially high bandwidth usage and unavailability of control functions during communications and cloud outages. This can be a problem in industrial applications that may require fast response times and a certain degree of local autonomy, including the ability to provide at least the basic control at all times regardless of cloud availability. Placing of processing and storage functions closer to the edge tends to alleviate these problems.

Data-processing functions can take place at more powerful gateways or at some variants of fog nodes that can perform some analysis and storage functions, as illustrated in Fig 1.3. In practice, the term fog tends to be used for relatively higher-

powered edge nodes that are used for processing and storage and may or may not provide communications gateway functions and interface to sensors. To illustrate this point, Fig 1.3 shows one potential placement of a fog node somewhere within the communications layer and possibly touching or residing in the cloud. Fog nodes are supposed to bring cloud-style processing and considerable power closer to the edge. The name is a metaphoric allusion to fog as the part of cloud that is close to the ground.

Rather than conceptualizing processing and storage nodes as being of a given type, placement of data-processing in IoT systems should be viewed as a functional continuum that can span the range from edge nodes all the way to the cloud. As discussed in subsequent chapters, IoT systems should be implemented to allow flexible placement and late mapping of functions to hardware. This allows allocation of functions to meet the needs of a particular system and to even respond to load changes dynamically by reallocating them at run time if needed.

Communications Layer

The communications layer enables a vast array of edge devices and things to exchange messages with each other, the rest of the IoT system, and ultimately the Internet. In IoT systems, the communications layer may include a variety of wireless and wired links, spanning local areas and including long-haul connections, such as wide-area networks and telco LTE variants. It may represent a complex infrastructure of links, bridges and routers that can transport payloads from local point-to-point segments all the way to any endpoint and application on the Internet. In addition to hardware, this requires implementation of a number of network layers and protocols which in IoT systems are commonly based on the Internet blueprint as discussed in Chapter 3 on communications. Moreover, whenever possible IoT nodes should make use of the already installed Internet infrastructure with its global reach and numerous access points.

The choice of a specific type of IoT physical connectivity can be a complex decision involving considerations such as system topology, range limits of various wireless technologies, location of usable Internet access points, bandwidth requirements, and costs of communication links in terms of initial interfaces and setup as well as the ongoing operational and bandwidth expenses thereafter. Another consideration in wired system is the cost of acquiring and laying out wires from sensors to gateways, which can influence the number and placement of gateways in the field.

Cloud

The cloud is the ultimate meeting place of large-scale IoT data, applications, and services that operate on them. Streaming and archived data are made available to

authorized applications and services via APIs and queries. Cloud implementations generally provide the top-end data aggregations, and the server, storage and connectivity infrastructure to execute services and applications that use them.

In terms of implementation, the term cloud generally refers to the large collections of servers and storage located in data centers that can be accessed via the Internet and allocated to services and applications in a dynamic and elastic manner. The details of cloud operation in IoT systems are discussed in Chapter 4

IoT data coming to the cloud may be processed in-flight as they arrive, stored for subsequent processing, or both. System-level rules and policies in effect determine how individual data should be processed and routed to the appropriate destination points and services.

The cloud portion of Fig 1.3 also illustrates the logical functions of aggregation and storage of sensor data and metadata. As discussed later, in the actual implementations they may be kept as separate or combined into a single database.

Analytics and ML algorithms may be applied to create insights and actionable recommendations by processing data at various points in an IoT system. As discussed in Chapter 4, such algorithms tend to be developed in the cloud where they can use the massive processing resources and large data aggregations, both of which are required for algorithm training and producing of system-wide insights. When completed, ML algorithms can be deployed at other nodes in the system, such as the edge, to perform inferences and spot trends using locally available data. Figure 1.3 also shows a separate storage for the analytics, which should be regarded as a logical function indicating the need for storing of derived data and algorithms.

An optional data-related function shown in the cloud in Fig. 1.3 is the interface to enterprise systems and online transaction processing systems (OLTPS). In various installations, enterprise systems may include any combination of systems such as ERP (enterprise resource planning), CRM (customer relationship management) and variants of e-commerce. In general, is useful to interface them to the IoT portion for data exchanges and to facilitate holistic insights that can benefit all sides of operation of a complex system.

Figure 1.3 also highlights two important interface points in an IoT system, notably the data and metadata APIs between the edge components and the cloud, and the data-retrieval APIs for live streams and stored data used by the applications and services executing in the cloud. In general, APIs should enable applications to query, search, and access data and metadata of interest. Formalizing the types of interactions and data formats that they support is not only a good design practice, it also provides the foundation for modularity and interoperability in implementation of IoT systems.

Control Plane

The components and functions described so far implement the primary objective of an IoT system to collect, process, and act on data. Parts of the system that carry

out these production activities and implement the related system flows are commonly referred to as the *data plane* or user plane. The task of keeping the IoT infrastructure itself running and secure is usually delegated to a sperate system overlay that is referred to as the *control plane*. It is partially depicted in Fig. 1.3 as the service and database labeled Security, Management and Provisioning. Although seemingly playing a supporting role to the primary mission, security and management are essential to keeping an IoT system up and running securely and with integrity so that it can fulfill its intended purpose and not be a threat to safety of the people and the environment.

During normal operation, control-plane systems constantly monitor activities that may impact security and availability. This is accomplished through the network of management agents that are installed on nodes and system components to observe and report status and changes. Their reports are customarily aggregated and visualized to operators at a central control point. The agents are also used to distribute and manage security policies and credentials, change configuration, and update firmware and software as necessary.

An important function of central monitoring is the detection and analysis of suspicious behaviors that may indicate probes from the adversaries or breaches of security. When incidents are detected, the handling mechanisms and policies are activated to mitigate the situation by identifying and isolating the compromised parts of the system. Details of those operations are presented in Chapter 5.

While the system is in operation, new nodes may have to be added, existing ones patched, and old ones decommissioned without bringing down other parts of the system or relaxing its security posture. Security and management systems are involved in preparing nodes for joining the system in the early stages of their lifecycle that precede activation. During the process of node commissioning and provisioning that follows installation, they are issued system identities and security credentials necessary for authentication and for secure operation upon their activation in the system. During that time, nodes are also entered into device registries and other backend systems that may need to be involved in their operation, such as the billing, asset management and support.

Book Organization

The rest of this book is organized as follows. Chapters 2 through 5 cover in detail the concepts and design considerations of system elements and components briefly overviewed in this chapter – edge, communications layer, cloud, security and management.

Chapter 2 covers the edge, starting with sensor data acquisition and processing, and continuing to the edge functionality that includes event processing, storage, local control and scripting, and interfacing to sensors and actuators as well as to the external communications and the cloud. After discussing the tradeoffs involved in

the functional placement of components in distributed systems on the edge-to-cloud continuum, including the fog, the chapter continues with a description of hardware and software considerations involved in the edge node design. It concludes with a brief description of the architecture and implementation of an open-source edge framework as a practical instantiation of the concepts covered earlier in the chapter.

Chapter 3 focuses on communications. It describes a layered network design which is the underpinning of the Internet and a useful blueprint for the IoT system design. It continues with the coverage of wireless and constrained networks at the edge, including the IEEE 802.15.4 and its derivatives and variants, followed by the description of 6 LoWPAN bridging to the Internet, and some non-IP based networks that have a large base of installed devices, such as the Zigbee and Bluetooth. Subsequent sections cover cellular offerings in the licensed spectrum, including their IoT adaptations, such as the NB-IoT. The chapter concludes with the exposition of the Constrained Application Protocol (CoAP) which is commonly used in IoT systems as a lighter-weight functional substitute for the HTTP, and the popular messaging and queuing implementation of the IoT publish-subscribe mechanism MQTT.

The Chapter 4 focus is on the cloud. It covers key elements and functions of IoT cloud core components, including data ingestion via edge-cloud gateways, in-flight stream processing, and short-term and long-term storage systems suitable for IoT applications. The second half of the chapter focuses on analytics and optimization algorithms. It starts with a real-life example and side-by-side comparison of the effectiveness and use of the algorithmic machine analytics and the traditional optimization methods. The rest of the chapter covers principles of machine learning, operation of artificial neurons and networks, and types and uses of ML systems. The closing section discusses the process and details of creating and operating a ML model, including model selection, training and optimization.

Chapter 5 covers the control plane, security and management systems. It covers types of security threats and attacks in IoT and OT systems, followed by the security planning and analysis steps that include risk assessment and threat modeling that indicate what should be guarded against and which mitigation techniques to use. A section on cryptography overviews key foundational elements of security design, including symmetric and public-key cryptography, key exchange, message authentication, and digital signatures. This is followed by the treatment of endpoint security, including hardware security modules, such as TPM and TEE that facilitate secure booting and software execution, followed by the software isolation mechanisms, such as virtualization and containers. The section on network security covers transport-level security and network isolation and segmentation techniques. This is followed by a treatment of security monitoring, incident handling, and systems management in major stages of node lifecycle, including provisioning and activation. The last two sections cover privacy and a summary putting it all together section.

Chapter 6 focuses on IoT data formats and representation. It explains why IoT data representation needs to include semantic annotation which makes it different

from the primarily textual worldwide web. The common structure of IoT information models is described, including object types, attributes, interactions and links. A separate section covers IoT interoperability among nodes through the use of shared information models. This is followed by the description of data serialization and format of payloads as they are exchanged on the wire. A section on metadata covers its definition, importance and provides examples of use and benefits. IoT frameworks are introduced as one approach to achieving functional interoperability among compliant nodes within a domain. Cross-domain interoperability is covered in a separate section. It identifies three levels of interoperability and focuses on the value and need for interoperability across domains and in large data aggregations that are a prerequisite for holistic system management and AI.

Chapter 7 covers IoT data standards and highlights the salient features of several of them with a focus on data information models and interoperability. It illustrates the scope and directions of the ongoing standardization work and reviews some of the more influential efforts including IPSO, OCF, WoT, Haystack, and OPC UA. Examples of their definitions of a basic sensor type are included for comparison. The chapter concludes with a summary of the commonalities, differences and some limitations of the presented approaches to the problem.

Chapter 8 provides an overview of several major commercial IoT platforms to illustrate the scope of what is available to IoT system designers as potential building blocks. It also points out the structural similarities in the architectures of the presented systems, and differences in their scope and emphasis.

Chapter 9 summarizes design and integration considerations involved in putting together an entire IoT system. It is intended to serve as a high-level checklist and an expansion of some of the issues and recommendations provided in the prior chapters on system components. It also includes a detailed example of the design of an actual system, starting with a definition of purpose and specification and design outline, followed by the implementation, experimental results, and a pilot evaluation with the actual users. It outlines key stages in the system development from the inception to completion and includes a discussion of redirections and changes made based on the insights gained in the design and implementation process.

References

[1] Gabbai, A. (2015) 'Kevin Ashton describes the internet of things' *Smithsonian Magazine* [Online] Available at: https://www.smithsonianmag.com/innovation/kevin-ashton-describes-the-internet-of-things-180953749/ (Accessed Dec 15, 2019)

[2] McKinsey Global Institute (2015) 'The internet of things: mapping the value beyond the hype' [Online] Available at: https://www.mckinsey.com/~/media/McKinsey/Industries/Technology%20Media%20and%20Telecommunications/High%20Tech/Our%20Insights/The%20Internet%20of%20Things%20The%20value%20of%20digitizing%20the%20physical%20world/The-Internet-of-things-Mapping-the-value-beyond-the-hype.ashx (Accessed Dec 15, 2019)

[3] ITU-T Y.4000/Y.2060 (2012) 'Overview of the internet of things' [Online] Available at: https://www.itu.int/ITU-T/recommendations/rec.aspx?rec=y.2060 (Accessed Dec 15, 2019)
[4] Wikipedia 'Internet of things' [Online] Available at: https://en.wikipedia.org/wiki/Internet_of_things (Accessed Nov15, 2019)
[5] Maker.io [Online] Available at: https://www.digikey.com/en/maker/ (Accessed Nov 15, 2019)
[6] Makezine [Online] Available at: https://makezine.com/ (Accessed Dec 15, 2019)
[7] Raspberry Pi [Online] Available at: https://www.raspberrypi.org/ (Accessed Nov 15, 2019)
[8] Arduino [Online] Available at: https://www.arduino.cc/ (Accessed Dec 15, 2019)
[9] SeeedStudio [Online] Available at: https://www.seeedstudio.com/ (Accessed Nov 15, 2019)
[10] Hanes, D or Barton P.? et al. (2017) *IoT fundamentals: networking, technologies, protocols and use cases for the internet of things* Cisco Press, Indianapolis, IN.
[11] Chou, T. (2016) *Precision: principles, practices and solutions for the internet of things.* Cloudbook Publishing, USA.