4

Cloud and Fog Computing in the Internet of Things

Daniel Happ

Telecommunication Networks Group (TKN), Technische Universität Berlin, Berlin, Germany

4.1 Introduction

In contrast to classical wireless sensor networks (WSN) that usually only serve a single application, one of the core benefits of the shift to the IoT lies in the common usage of sensor hardware by heterogeneous applications (Tschofenig et al., 2015). Additionally, the revolution of the IoT does not stem from the number of connected things alone, but from the solutions and services offered on top of the data. The basic requirements of such value-added services can be briefly summarized into nonvolatile storage of historical sensor data, sensor data processing, and efficient near-real-time distribution of sensor data.

However, although everyday objects are increasingly connected to the Internet and becoming more and more powerful, they are usually not capable enough to fulfill all those requirements themselves. One of the major challenges is that commonly used devices, such as sensor nodes, smartphones, and wearable techs, usually run on battery power, making storage or complex processing of a large amount of data unfeasible. Mains-operated connected objects may also often be too constrained to perform those tasks as reliably and quickly as required.

The recent advances in cloud computing have led to an increasing usage of this model to meet the aforementioned requirements and enable value-added services in the IoT context. Cloud computing offers convenient, ubiquitous, and on-demand access to a shared pool of configurable computing resources, which are accessible over the Internet, and usually reside in third-party datacenters (Hassan et al., 2012; Mell and Grance, 2011). Along with these resources, cloud providers offer fast and configurable networking for data distribution and reliable, nonvolatile, replicated storage. Thanks to its flexibility, reliability, and

Internet of Things A to Z: Technologies and Applications, First Edition. Edited by Qusay F. Hassan.

usage-based cost model, cloud computing is well positioned to meet the specific requirements of value-added services in the IoT context. Constraint devices can save energy by transferring their data to a cloud-based platform where it will be distributed to multiple relevant applications and services, which will process the data accordingly.

Though this architecture works well today, it is not suitable for latency-sensitive applications since cloud datacenters are colocated neither with connected objects nor with consumers of value-added services (Zhang et al., 2015). From a network topology view, cloud datacenters are located several hops away from IoT data producers and consumers, and are most often separated by a constrained last-mile link. The physical distance alone causes additional latencies that may not be acceptable for latency-sensitive applications such as control loops. Thus, instead of forcing all IoT communications through a cloud intermediary, there has been a push to move storage, processing, and distribution of data closer to the edge of the network, toward data producers and consumers.

One of those concepts is fog computing, which is seen as an extension of cloud computing, but enable services closer to the edge of the network (Bonomi et al., 2012). While fog and cloud computing paradigms share many mechanisms, fog computing mainly addresses applications and services that would not be feasible in the cloud. One reason would be to eliminate bandwidth bottlenecks and improve on latency for local control loops commonly found in the IoT context. To achieve this, fog computing uses processing power that is available locally today, such as on network hardware or local gateway nodes, mobile phones, or additional hardware that would have to be deployed in the future between the user and the cloud.

This chapter gives an introduction to the concepts of cloud- and fog-based computing in the IoT context. For both approaches, advantages and challenges are presented. Additionally, specific IoT use cases are outlined that benefit from cloud and fog computing, respectively. Since fog computing is a rather new model, the overview also includes potential future use cases that are not yet fully realized.

4.2 IoT System Requirements

A central concern when analyzing cloud and fog computing is the actual role these technologies are expected to play in IoT systems. As discussed in Chapter 1, the basic architectural reference model has three layers: device layer, connectivity layer, and application layer. In general, cloud and fog technology will be used to realize the connectivity layer by relaying device data and enabling value-added services in the application layer. In particular, this chapter identifies the following basic functions as what the cloud or fog backend must provide for IoT systems:

1) *Data Distribution*. IoT environments will generate enormous amounts of data that can be useful in many ways. The first requirement is to provide ubiquitous, cross-vendor real-time access to the data by distributing it from data sources (e.g., sensors, wearables, and smartphones) to consumers (e.g., actuators, value-added services, and applications). The great diversity in hardware and software highlights the need for a unified messaging middleware interconnecting data sources to consumers by providing uniform and standardized APIs.
2) *Scalable Storage*. The value of the vast amount of data available from things might not be directly evident on collection, but the data might provide very valuable insights in the future. Hence, for many use cases, scalable storage of and ubiquitous access to historic data will be a valuable function.
3) *Processing Services*. As mentioned in Chapter 1, analytics and big data processing have become increasingly relevant to the IoT ecosystem. Raw IoT data alone are not of much value, but value-added services can use analytics as well as temporal and spatial aggregation and correlation to further provide important insights.

Apart from those basic services, there are additional high-level qualities imposed by the specific IoT use case:

1) *Flexible Self-Organization.* As stated in Chapter 1, the IoT has to be self-organizing. While things are to be uniquely identifiable, new data producers may join or leave the system at any time, so it must enable the discovery of relevant data sources and services out of a heterogeneous and constantly changing blend. On the other side, the system has to be able to automatically adapt to the changing needs of data consumers.
2) *Reliability.* Depending on the particular use case, stringent reliability and quality of service requirements must be respected, such as low latency and reliable end-to-end data delivery.
3) *Scalability.* Due to the large number of connected things as well as data consumers that are expected, the system has to be scalable, that is, the system should provide adequate basic services independent of the number of connected devices, services, and consumers.
4) *Data Confidentiality and Security.* Device data or derived insights might be sensitive, possibly also not allowed to be shared with certain entities or not allowed to be transferred or stored in other juridical areas. The IoT system has to support those requirements and ensure data are not accessed by restricted entities.

The aforementioned requirements motivate the use of cloud and fog computing to enable the IoT vision. The chapter continues by giving a brief introduction

to cloud computing and how it is used in IoT systems; then fog computing and its integration in modern IoT systems is also addressed.

4.3 Cloud Computing in IoT

Cloud computing refers to both a subset of applications delivered as services over the Internet and the underlying hardware and software systems in the datacenters that provide those services. Cloud computing is usually divided into Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS), as shown in Figure 4.1 (Hassan et al., 2012; Mell and Grance, 2011).

IaaS describes a business model in which no complete solution is offered, but only the hardware necessary to implement specific applications. PaaS, on the other hand, allows developers to develop and run their applications on the managed infrastructure using well-defined interfaces. For this purpose, PaaS providers usually maintain development environments in the form of frameworks. SaaS is a model in which the software application is no longer sold to the customer for a one-time fee, but is provided as a service for a recurring fee. In this model, infrastructure and platform are also fully managed by the service provider.

Cloud computing can help realize IoT systems that meet most of the requirements mentioned in the previous section. Indeed, the approach to send device data to a cloud-based service for messaging and processing is widely adopted by developers (Gubbi et al., 2013; Menzel et al., 2014). Additionally, solutions for many common tasks are already offered as hosted services from cloud providers, including storage, messaging, and processing. For instance, there are PaaS providers specifically for IoT systems, such as Amazon IoT, that offer a so-called device gateway for IoT messaging needs, a rule engine

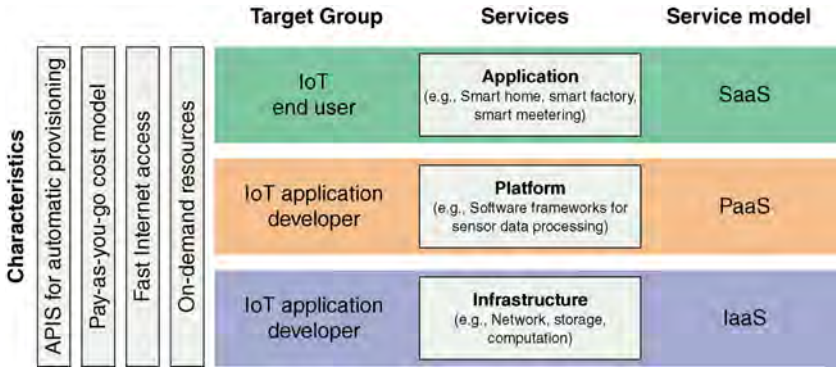

Figure 4.1 Cloud computing service models.

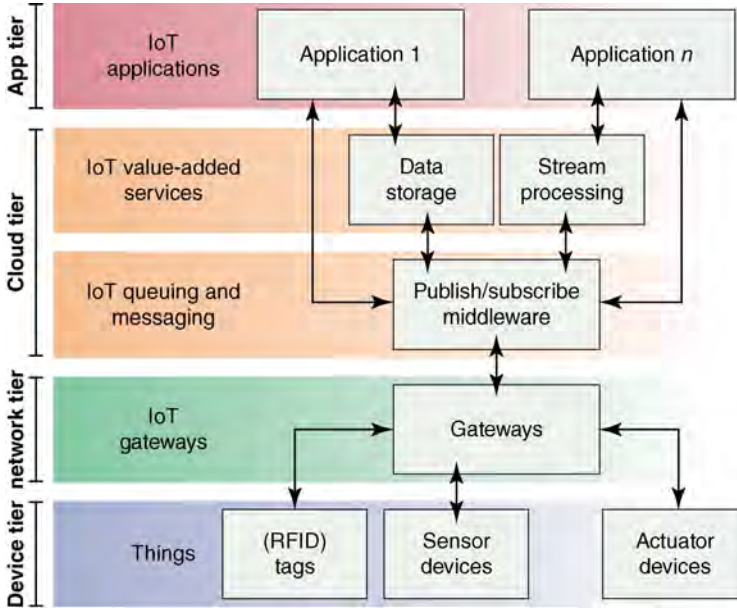

Figure 4.2 Common cloud-based IoT architecture.

for data processing, and a database backend (Dynamo DB) for storage and querying. Other vendors, such as Microsoft, IBM, and Xively, also offer similar services. There are also SaaS providers in the IoT ecosystem that offer ready-made services for specific purposes, such as smart heating or building automation. Examples include Nest[1] and Tado,[2] which market smart thermostats for home users, aiming to reduce the users' heating cost.

Most IoT service providers use a three-layered architecture similar to the one depicted in Figure 4.2. The typical cloud-based IoT solution is characterized by many data-producing end devices that usually use gateway nodes to communicate to a wide area network, usually the public Internet. The storage, messaging, and processing that are crucial for IoT applications are offered as services in the cloud.

An overview of the general placement of things, gateways, and cloud is shown in Figure 4.3. Usually connected things themselves are severely constrained and connect to a gateway, which is connected to the Internet. However, most often this link uses a relatively slow last-mile technology, such as DSL (digital subscriber line) 3G/4G, making it the bottleneck between devices and the cloud in terms of delay and throughput. Regarding scale, because more resources are

1 https://nest.com
2 https://www.tado.com

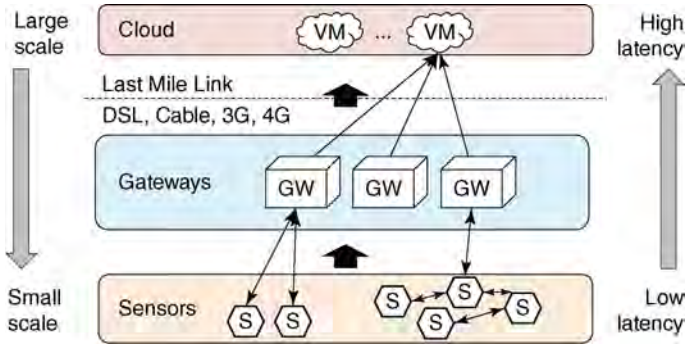

Figure 4.3 Scale and latency in cloud-based IoT systems.

available, the number of supported services and the potential number and spatial distribution of data producers and consumers increase toward the cloud. However, the latency from the end-devices to storage and service nodes also increases toward the top.

The queuing and messaging layer often uses an event-driven publish/subscribe system in the cloud-based IoT approach (Happ et al., 2015; Antonic et al., 2014; Hunkeler et al., 2008; Banks and Gupta, 2014). As a result, data producers and consumers are loosely coupled in space (do not have to know each other), time (do not have to be available at the same time), and in synchronization (asynchronous, nonblocking messaging, that is, calling applications do not have to wait for system to complete call).

4.3.1 Advantages of Using the Cloud for IoT

Cloud computing offers the means to meet some of the IoT systems' requirements outlined above.

1) *Flexibility*. Cloud instances are usually based on virtualized physical hardware in datacenters. Resources are pooled and can be used in an on-demand fashion, often with a pay-as-you-go pricing model, where only resources actually used are billed. This enables on-demand resource provisioning matching the constantly changing requirements of IoT applications.
2) *Reliability*. Intra- and inter-datacenter redundancy can be used to create scalable and reliable services, such as distributed object storage. The basic services of distribution, storage, and processing can be easily realized using a cloud-based infrastructure.
3) *Fast Ubiquitous Access*. Cloud datacenters are well connected using both private peering and multiple uplink transit providers. This enables ubiquitous and fast access to services from anywhere over the Internet.

4) *Fast Time to Market.* Cloud providers usually offer readily available software components to use for the development of customized solutions, which can be accessed through well-defined APIs. Some providers specifically target IoT use cases. This reduces both the time to market for applications and the development cost.

4.3.2 Examples of Cloud-Based IoT

To further emphasize the impact of cloud computing on the current IoT ecosystem, the chapter continues with an overview of recent use cases where the use of the cloud-based IoT paradigm was proposed. The following examples are based on the taxonomy of IoT domains that was introduced in Chapter 1, namely, industrial, smart city, and health care domains.

4.3.2.1 Industrial Domain

The combination of IoT and cloud computing makes several innovative use cases in the industrial domain possible (Tao et al., 2014). Smart industrial IoT systems allow monitoring of industrial plants to not only improve productivity but also safety, especially for high environmental risks. For instance, the approach has been used for a monitoring and prealarm system for tracking tailings dam failures in the mining industry (Sun et al., 2012). The system offers real-time monitoring of the saturated line, water level, and dam deformation. Data are gathered by a local sensor network and processed in the cloud. The results are disseminated to Personal Digital Assistants (PDAs) and other client devices. The system has since been applied and proven useful in several mines.

The combination of IoT and cloud computing can enable advanced solutions in the transportation industry as well. Sensors in vehicles can monitor various aspects related to the efficiency of the transport: tire pressure, fuel consumption, location, and speed (Borgia, 2014).

4.3.2.2 Smart Cities

IoT together with cloud-based analytics has the potential to increase the quality of everyday urban life. Particular focus lies in more efficient and environment-friendly use of energy, transportation, environmental protection, and urban planning.

Efficient management of renewable energy sources is a core building block for achieving sustainability in urban environments. The future smart grid will transform the electrical distribution infrastructure into an intelligent multi-directional system, which will enable the decentralized generation of energy and real-time transparency for producers and consumers. Application areas include demand response and forecasting, dynamic pricing, micro-grid management, and real-time monitoring (Bera et al., 2015).

One application of cloud computing in this context is the use of readily available cloud storage for data from smart meters and other related sensors. Cloud storage allows not only the necessary data storage but also easy access and means for additional analytics. It further provides high fault tolerance through redundancy and can use versioned copies for transparency and roll-back (Rusitschka et al., 2010; Fang et al., 2012).

The combination of cloud computing, IoT, and crowd sensing can further be used to protect the environment and raise awareness of inherent problems, such as the air quality. For example, a crowd-sensing application can be used for monitoring urban air quality, using sensors that are carried by citizens across the city (Antonić et al., 2014a, 2014b). The data gathered are sent to cloud servers, which will filter out the redundant or unneeded data. This particular system additionally provides personalized real-time notifications about air quality to mobile users using Google Cloud Messaging.

SmartSantander (Sanchez et al., 2014) provides an experimental research facility in a real urban deployment in the smart city context. It focuses on the sensor network part of smart city research and provides a large-scale sensor testbed that help develop solutions for connecting and managing smart city sensor devices. To support applications on top of the data, a private cloud infrastructure is provided.

4.3.2.3 Health/Well-Being

IoT will enable applications in the smart health care and well-being domain that will improve the experience for patients and physicians alike. The spectrum of useful applications ranges from remote health monitoring to drug inventory management. For various use cases, offloading heavy processing to cloud-based backends has been proposed.

A key challenge for health-related use cases is the remote monitoring of patients (Delmastro 2012; Hossain et al., 2016; Kan et al., 2015). There have been multiple proposals to bridge the gap between physical world and software using cloud-based approaches. Medical parameters and vital functions (e.g., blood pressure and heart rate) can be monitored in real-time and data gathered using Body Area Networks (BANs), which pass the data on to gateways and ultimately to a cloud backend. For instance, data from an electroencephalogram (EEG) can be collected using connected devices and sent to cloud-based virtual machines, which will further analyze them (Peddi et al., 2017).

Another important aspect in this area is the well-being of users. Smart applications, such as fitness trackers, can help motivate its users to exercise on a regular basis by giving positive feedbacks, such as the number of steps walked during a day. Many applications can make use of knowledge about the patients' mood or feelings. Multiple readings from sensors already deployed can be combined for emotion detection. For example, Hossain et al. (2016) use audio

and visual data from a camera that might be available in future smart homes, a cloud-based Hadoop cluster, and 5G networking to achieve over 80% accuracy in detecting emotions. They also show the reduction of processing time when using multiple cluster nodes in contrast to running the same workload on a single server. The presented solution thus is an example of how the on-demand resources available in the cloud can be used to speed up IoT-related processing services, especially for data-intensive workloads such as video. Similar work has studied a feature extraction method for emotion recognition using EEG-based human brain signals (Mehmood and Lee 2016).

4.3.3 Key Challenges of Cloud-Based IoT

There are also several specific use cases where cloud computing is not prepared for IoT applications. The limitations of cloud computing in the IoT context can be summarized as follows:

1) *High or Unpredictable Latency.* Low and predictable latency are fundamental requirements for many control loops, for instance, those commonly found in industrial automation, but also other IoT use cases, such as home automation. As cloud datacenters are mostly located where energy is inexpensive, they are not located close to potential IoT users, adding unavoidable latency to control loops, since the physical distance between users and the cloud dictates the minimum latency that can be achieved.
2) *High Uplink Bandwidth Requirements*. Gateways that do not have the bandwidth capacity to upload certain types of sensor data to the cloud will not be able to use the cloud-based storage and processing approach. For instance, a video analysis value-added service may not be feasible as a remote cloud-based service, since the data that have to be transferred are too large for the link available. This is especially true for rural areas or mobile settings, such as a 3G uplink.
3) *No In-Network Filtering or Aggregation.* Some applications cover a large geographical space, whereas only an aggregate value of the sensors is actually important. For instance, if the maximum temperature across various locations is important, sending every sensor reading to the cloud will not be the most efficient solution. Instead, in-network processing could greatly reduce the amount of data that actually need to be sent to, processed on, and stored in the cloud.
4) *Uninterrupted Internet Connection Required.* Some applications, such as smart connected vehicles, may not have a connection to the Internet and, thus, to the cloud at all times. During network outages, the cloud approach would stop working, and in such cases, local actuation loops would not run. More local processing might be used instead to at least provide a fallback provisional service until network connectivity is restored.

5) *Privacy and Security Concerns.* Cloud providers are usually expected to respect the privacy of the user and provide sufficiently secure services. However, there is no easy way to measure or monitor the security of cloud services, so ultimately those requirements cannot be easily verified and the provider has to be trusted to some extent. Even if the provider is trusted, since resources are virtualized, containers of different tenants are located on the same physical machine, where software bugs could potentially leak private data to third parties. Moreover, laws might demand that certain data must not be stored outside certain juridical areas, which are not obviously verifiable with most cloud providers.

These limitations imposed by the use of cloud computing motivate the need for an additional technology that can mitigate these issues especially in critical IoT applications. Fog computing, as will be discussed in the following sections, seems like a viable option in such an environment.

4.4 Fog Computing in IoT

Although the benefits of cloud computing in the context of IoT are widely recognized by industry and research communities, there has been criticism about offloading all data to and processing it on the cloud (Happ and Wolisz 2017; Zhang et al., 2015). In many use cases, the consumer of data or the user of related services is in the vicinity of the data producer, but still the data take the route via the cloud. One problem with this is the unnecessary load on the ISPs (Internet service providers) and the additional delay introduced. The unnecessary traffic could make the participation in the network expensive or even prohibitive where no broadband connection is available and alternative technologies, for example, cellular data links, have to be used. This concerns both developing countries and rural areas where sensors have to be deployed.

Fog computing is a term coined by Cisco for a concept similar to cloud computing that offers a highly virtualized resource pool at the edge of the network (Bonomi et al., 2012). It provides computation, storage, and networking services to nearby end users, opposed to the cloud that is typically located at the edge of network. Hence, the fog will have a special widely distributed deployment, in addition to a large heterogeneity of devices. The fog concept can be illustrated as a cloud near the ground or the user. Traditional content delivery networks (CDNs) share a similar concept of bringing data closer to the user, which helps in meeting the increasing demand of streaming media traffic by placing data at the edge of the network. Akamai's Edge Computing, Cisco's IOx, and Intel's Intelligent Edge extend this concept by bringing cloud-like services—in particular, the execution of customized value-added services—closer to the user.

Microsoft has promoted a similar concept with their micro-datacenter approach as a smaller version of the cloud (Brown 2015). The envisioned micro-datacenters are self-contained computing environments with computing and storage resources, and they are connected to the Internet using high-speed connections. They basically follow the typical cloud approach, but bring the hardware to the premise on a smaller scale to avoid unnecessary delays. A customer may host tens of servers on the premise that have single-digit terabytes of storage when combined.

Fog computing has substantial overlap with edge computing, which is a similar technology that enables storage and processing at the edge of the network. The research community has not yet come to a precise definition of fog computing and the specific differences to edge computing. However, one work (Garcia Lopez et al., 2015) mentions human-driven applications, such as video streaming or web browsing, which are usually triggered by human interaction, as an indicator for more traditional edge computing, while machine-to-machine communication as seen in the IoT would be an indicator for fog computing.

The overall architecture including cloud and fog instances is depicted in Figure 4.4. Instead of a central cloud layer, several hierarchical layers of fog nodes are introduced, which are increasingly closer to the edge of the network. Gateways are now seen as part of an abstract fog layer. Since those devices are

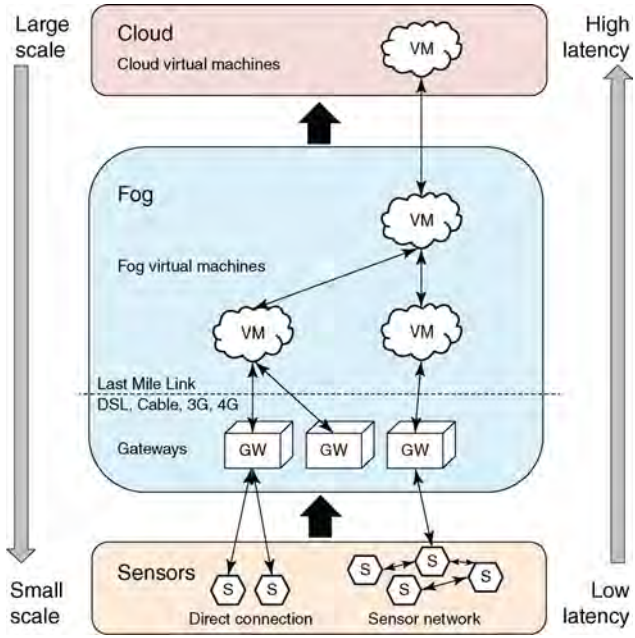

Figure 4.4 Revised fog-enabled architecture combining fog and cloud-based IoT.

Table 4.1 Overview of some gateway devices suitable for local computation, processing, and storage.

Device	CPU clock	Memory	Storage	Price
Intel NUC	1.3 GHz	16 GB	SATA	~$300
BeagleBone Black	1 GHz	512 MB	4 GB EMMC	$55
Raspberry Pi 3	1.2 GHz	1 GB	SD	$35
TP-Link TL-WR841N	650 MHz	32 MB	4 MB	420

still constrained in terms of processing power available, the remote cloud provides storage and processing capabilities when they are not sufficiently available in the fog. Generally, fog instances will also be available on other network hardware, such as the routers at the ISPs, providing the means to analyze and process data within the network closer to the end user than in a centralized remote cloud. More powerful fog nodes may also offer the provisioning of virtual machines (Yi et al., 2015; Osanaiye et al., 2017). However, while available resources and thus scalability increase toward the top, the physical distance, latency, and cost of communication also increase.

Local processing on gateway devices is made possible by advanced IoT hardware, which includes powerful smartphones and embedded single-board computers. While most IoT devices are still energy constrained, they will mostly use a gateway to connect to the Internet. Those gateways, on the other hand, are usually mains-powered and powerful enough to take over some of the tasks that cloud services are providing today, such as data processing and storage. A selection of possible gateway devices that could be used today to do the interfacing to low-power sensors and additionally offer local processing and storage is compiled in Table 4.1. The table shows that some of the devices available today potentially have a lot of processing power and storage available that is probably not fully utilized all the time. Gateway hardware could be a good option for data processing and storage since the delay-and-bandwidth-constrained Internet uplink would not have to be used at all for local processing loops.

4.4.1 Advantages of Using the Fog for IoT

This section highlights the advantages of extending cloud computing with closer-to-the-edge fog computing resources. There are several reasons why utilizing processing power that is available at the edge is useful:

1) *Minimizing Latency.* Some use cases are basic control loop; that is, on a certain condition, a certain action is triggered. These control loops, however,

are often quite time-sensitive or users expect a certain action almost instantly. Since cloud datacenters are usually deployed where it makes sense economically, the physical distance between the data source and the utilized cloud service results in a concerning latency. Analyzing raw data and making decisions locally or close to the data source can greatly reduce the latency of those control loops, since the distant cloud service will not be contacted.

2) *Improving Reliability*. The vision of IoT includes using sensor data for public safety or critical infrastructure. While most enterprise cloud providers offer very reliable datacenters with redundant network connections, storage, and processing infrastructure, the uplink to the distant cloud could turn out to be easily breakable in those scenarios. It would be very valuable to be able to at least do some of the processing locally as a fallback option, or exclusively use local processing altogether.
3) *Addressing Privacy Concerns*. Some IoT data will be sensitive or required by law to not be stored outside specific geographical boundaries. While cloud service providers are usually considered trusted, the user has no control over where the data are actually stored and who can access it. It also cannot be fully ruled out that due to bugs, third parties could potentially unlawfully access sensitive data. Local gateway or other edge nodes, on the other hand, can be trusted since they are under the control of the local operator. Of course, this is assuming that no security vulnerabilities in the local systems.
4) *Conserving Bandwidth*. The uplink bandwidth of IoT gateways is often severely limited, such as DSL or a 3G connection. It is not always feasible to transport vast amounts of data from edge devices to the cloud. Performing data processing locally and only sending aggregates and filtered data to the cloud can significantly reduce the uplink bandwidth needed, making it possible to implement IoT systems that are connected to the cloud using limited or intermittent connections.

4.4.2 Potential Future Fog Use Cases in the IoT

There are various possible use cases where IoT applications can take advantage of fog computing. The following sections provide some examples of such applications.

4.4.2.1 Smart Grid

Smart grid is the next-generation electricity network that aims to provide more effective load balancing and higher reliability, as well as reduced electricity cost by automating the metering process. Smart meters will collect and relay information on electricity consumption, enabling more fine-grained pricing models based on actual supply and demand. Fog computing can be used for storing metering information locally in such scenarios, where local gateways would act as the lowest tier of semipermanent storage. Fog nodes between end

user and the cloud would only be updated in batches, reducing bandwidth requirement for smart meters. Instead of relying on higher layers, such as the cloud tier, for advanced analytics, fog nodes close to the user could use such information to either help customers change their behavior regarding power usage or to better predict future demand and supply. Fog or edge computing could also be used to improve local energy load balancing applications. Based on current price and local energy needs, the system can switch on or off appliances that need a lot of power, or it could switch them automatically to different power sources, such as solar and wind.

4.4.2.2 Connected Vehicles

Connected vehicles refer to techniques providing wireless connectivity for vehicles enabling direct communication between vehicles themselves, and between vehicles and the environment. The application that would benefit the most from fog-based approaches in that context would be the automatic intelligent reaction to sensor readings found in today's cars. For example, a smart traffic light could stop or slow down approaching traffic to avoid traffic jams or accidents when multiple braking vehicles are detected. Those use cases need inter-vehicle communication as well as local processing, as sending all sensor readings to a distant cloud is expensive and introduces a delay that is not acceptable for timely reaction to traffic conditions. Local storage and sensor data processing can help overcome those issues and enable those use cases.

4.4.2.3 Education

Students are increasingly using devices such as desktop computers, laptops, or tablets for their studies. This enables students to study at their own pace and to have the same information available at their homes as in the classroom. By tracking students' progress on these devices, teachers can gather real-time performance data on individual students and gain actionable insights into which students need further assistance. Analyzing performance data and storing additional course materials fundamentally pose the same challenges to the infrastructure as IoT data, which also has to be stored and processed. Fog techniques can also be used to enhance the privacy of student data by not uploading sensitive data about their performance to cloud-based systems.

4.4.2.4 Health Care

Sensors can collect health information, such as electrocardiograms, temperature, or blood glucose level. Additional environmental sensors could detect if elderly persons are following their daily routine, for example, they could monitor if the person gets up in the morning or is eating regularly. Most of these applications have high reliability and strict latency requirements. On top of that, the privacy of and regulations on patient data must be considered. Patient data could be directly transferred to physicians for diagnosis. Fog-based storage and

processing could enable the better monitoring of patients and elderly people without sacrificing privacy or reliability.

4.4.2.5 Smart Buildings

Smart buildings can optimize the energy consumption of buildings and increase comfort of living as well as safety. These smart environments are made possible by the fusion of multiple technologies, such as temperature and humidity sensors among other sensors, ubiquitous connectivity, and data analytics. Smart buildings could, for instance, detect when nobody is home and turn down heating. Gas or air quality sensors could improve safety by giving a warning signal for bad air quality or health-threatening concentrations of certain gases in the home or even take actions on their own by opening windows to let fresh air in. Of course, these actions can be triggered over a cloud-based service, but since data producers and consumers are mostly local, the increase in latency and traffic is unnecessary. A local fog-based system could use external data from cloud-based services, but analyze local sensors readings without sending them to distant clouds. As system reliability is crucial in some use cases, it is important to note that fog-based services would keep on working in case Internet connectivity is intermittent.

4.4.2.6 Surveillance

Surveillance, including intelligent video surveillance, relies on cloud-based systems for complex video analytics (VSaaS). Instead of a human operator, a cloud-based service employs computer vision algorithms and pattern recognition to evaluate the scene. Other sensor readings can help with the recognition of potential threats. A fog-based system can help improve the detection and response time of such a system by avoiding the delay introduced by communicating with a distant cloud service. Multimedia data in particular are large in size and it is not always feasible to send high-resolution high-frame-rate video to the cloud. Instead, the video can be analyzed locally, potentially at a higher resolution and frame rate than would be possible to send over the Internet uplink. While gateway hardware is constantly becoming more suitable for those computing tasks, they are still limited in their capabilities in comparison with the vast resources available in the cloud.

4.4.2.7 Wearables

The rise of wearable sensors, such as fitness trackers and smart watches, has led to an increasing demand in processing sensor data, such as activity recognition based on accelerometer and gyroscope readings. The data have to be sampled at relatively high frequencies to derive meaningful activities. The activities themselves might be highly private and need to be protected from unauthorized access. Finally, the result of the recognition algorithm may be needed with low latency in settings where no Internet connection is available. In these settings,

local processing on fog nodes might be a good compromise between offloading heavy computation from battery-powered devices and keeping data and processing close to the user.

4.4.2.8 Virtual Reality

In recent years, there have been several virtual reality solutions, such as Oculus Rift, HTC Vive, and Google Cardboard, that made it to the market and were inexpensive enough to see moderate adoption. In the future, some of these devices may become wireless and portable. The amount of data that has to be processed to make a convincing virtual reality experience possible could potentially be offloaded to fog nodes, saving energy and hence battery power, while still having a latency small enough to go unnoticed.

4.4.3 Examples of Fog-Based IoT

Recent research has highlighted the applicability of fog computing in the IoT context. A summary of the recent findings on the industrial domain, smart cities, and health is provided in the following sections.

4.4.3.1 Industrial Domain

In the industrial domain, there have been several proposals to process and store IoT data closer to the edge in order to overcome some of the challenges specific to the use case. For instance, a platform aiming to improve the effectiveness of applying equipment failure models was proposed (Gazis et al., 2015). The prototypical implementation uses existing sensors to monitor the deployed industrial machinery to detect potential anomalies. The platform builds failure models using machine learning algorithms and makes them available to the manufacturer and the administrator of the equipment. For moving this data processing closer to the edge, a Cisco ISR 819 router supporting the data in motion (DMo) framework is used. The study finds that fog computing in two realistic use cases reduces network traffic and increases resource utilization.

To support the messaging requirements of the industrial use case, an extension to the MQTT protocol was proposed (Peralta et al., 2017). A modified MQTT broker is placed at the fog layer, which supports predicting future measurements and the capability to offload computationally expensive data processing jobs from the cloud to the fog. The study shows that the energy efficiency is improved compared to using the standardized MQTT in those industrial settings using fog computing.

4.4.3.2 Smart Cities

Combining IoT with fog computing has the potential to move smart city analytics tasks closer to the user. Going in this direction, Stack4Things (Bruneo et al., 2016) was proposed, a model for fog-based IoT for smart cities based on

OpenStack, a software toolkit for controlling computing, storage, and networking resources in cloud datacenters. A smart mobility scenario in which a smart car interacts with smart city objects to achieve a variety of geolocalized services, such as signaling presence to traffic lights, is considered. A quantitative analysis shows that the framework allows managing of small areas within an urban region, locally on fog nodes.

Other works have focused on the analytics involved in the smart city domain. One work proposes a fog-based context-aware real-time data analytics platform (Jayaraman et al., 2014). The example use case of monitoring citizen activity in a smart city is given. The approach of combining local analytics, optionally combined with smart data reduction and on-demand sensing, is shown to significantly reduce the energy footprint of smart city analytics.

A recent study considers the use case of a smart pipeline monitoring system based on fiber optic sensors (Tang et al., 2015). It uses a hidden Markov model to sequentially learn and detect hazardous events threatening pipeline safety. Using a prototype and experimentation, it is shown that this is feasible even on a city-wide scale. Also, using the fog computing approach, the amount of data that has to be transmitted over the Internet, as well as the response time, is decreased substantially.

4.4.3.3 Health/Well-Being

Especially in the health domain, there have been privacy concerns when using cloud-based services. Fog computing could be one option to address these concerns by storing and processing data partially on devices the user trusts or even owns. A recent example is Health Fog (Mahmood et al., 2016), a framework that enables sharing and processing of health-related sensor data. The proposed system introduces an intermediate fog layer between the devices and the cloud to give the patient control over the flow of their health data and thus enable for better control over data privacy and security while still enabling processing and sharing of their data.

A design of a fog-based system to identify the chikungunya virus based on the patient's symptoms and surrounding environment conditions was also proposed (Sood and Mahajan 2017). It introduces different layers for different tasks and users. The task of the fog layer is to diagnose the virus and generate emergency alerts for patients and doctors. These alerts are relayed to a cloud-based service, which calculates the probability of the disease spreading and generates alerts for government agencies to counter the outbreak.

In a similar fashion, one work (Gia et al., 2015) studies ECG feature extraction on smart gateways. The study considers features that are important for the diagnosis of many cardiac diseases, including heart rate, P wave, and T wave. It shows that the approach is feasible by providing a proof-of-concept implementation on real hardware and show a reduction in traffic of more than 90% and a significant reduction in latency when using fog computing.

4.4.4 Key Challenges of Fog-Based IoT

Despite the advantages of fog computing, several issues need to be addressed toward the realization of the fog-based IoT approach as well as the integration with the current cloud-based model. Some of the key challenges are discussed in the following.

1) *Technological Interoperability*. The seamless interaction between devices and systems from different vendors is a major challenge for IoT as well as for the fog-based extensions to IoT. As discussed in several chapters of this book, there is still a lack of common standards for communication protocols, both locally and for the uplink to the cloud. For example, for connecting sensors to gateways, technologies used include Bluetooth Low Energy (BLE), 802.15.4 or ZigBee, or Wi-Fi (802.11). Also, long-range wireless technologies, such as LoRaWAN or Sigfox, are being used. Uplink protocols include proprietary technologies, in addition to open pub/sub protocols such as MQTT, XMPP (Extensible Messaging and Presence Protocol) or AMQP (Advanced Message Queuing Protocol), and request/response protocols such as COAP (Constrained Application Protocol). Protocols have to be standardized and agreed upon for the success of IoT. Additionally, techniques for storage and processing migration have to be developed.
2) *Semantic Interoperability*. For interoperability, it is also necessary that the devices and software involved interpret the information gathered, processed, and shared in the same way and act upon commands accordingly. A semantic model has to be available for every aspect of the fog/cloud approach; this should not only be for understanding the data but also for expressing the requirements and constraints when processing or moving these data, including quality of service or privacy requirements. There is still a lack of insight into which ontologies may be suitable for these tasks. Semantic web technologies could also help enhance device discovery or implement automatic reasoning.
3) *Programmability*. Data processing is of exceptional importance to the IoT ecosystem. Due to the variability of the requirements of event-based processing tasks, it is further crucial to automatically relocate processing tasks between the fog and cloud nodes. However, it is still unclear how processing jobs should be defined in the first place. There is still a lack of insight into which programming language and which interfaces are necessary for IoT data processing and for enabling seamless offloading of tasks between different systems, potentially using different hardware architectures and different formats. In particular, there has not been any consensus yet about whether to use an interface in the form of discrete functions or in the form of containers or virtual machine images.
4) *Scalability*. In the near future, the IoT will be composed of billions (or even trillions) of devices. The number of connected nodes will outnumber by

several orders of magnitude the number of hosts in today's Internet. Despite the problems that fog-based processing and storage may solve, it is still unclear how a seamless interaction between devices behind different gateways should work. Questions such as how devices should be discovered on a large scale, where the registry in a cloud–fog system should be located, and where data should sit to minimize communication latency and increase throughput are still not fully answered.

5) *Resilience and Reliability*. The fog-based approach is appealing to developers for implementing applications where relying on the permanent availability of the link to the cloud is not acceptable due to the possibility of temporary outages. Examples include industrial control loops or emergency response systems.

4.5 Conclusion

Cloud-based services are increasingly used to assist constrained IoT devices with storage and processing. The cloud offers various advantages in the IoT domain, including the flexible on-demand availability of resources, fast and reliable networking, and a multitude of hosted services ready to be used. While the benefits of cloud computing in the IoT context are evident by the vast adoption seen in existing systems, the main shortcoming is the position of cloud datacenters, which are often several hops away from data producers and consumers. Data usually have to pass a constraint link between data producers and the cloud and between the cloud and data consumers. This link adds additional delay and may limit the amount of data that can be transmitted. Over the next few years, fog computing will increasingly help to tackle those shortcomings by introducing cloud-like resources closer to the user.

This chapter gave an overview of the advantages, potential use cases, recent examples, and open challenges of cloud and fog computing in the IoT context. Both approaches have distinct benefits and will see wider adoption. Most likely, they will be combined to achieve a seamless integration of fog and cloud resources into a common IoT resource pool. Future developments will include mechanisms for fog–cloud interaction, such as automatic resource provisioning, replication, and migration, which are essential for meeting the IoT's requirements for resilience and reliability. Value-added service will have to be defined only once and will be automatically provisioned and relocated on demand to a suitable processing node. Existing services could be combined to create new innovative services. Advances in security and privacy will help keep sensitive data confidential and secure across processing nodes from different providers. Despite the remaining challenges, cloud as well as fog computing will prove to be indispensable tools in realizing this IoT vision.

References

Antonić, A. et al. (2014a) Urban crowd sensing demonstrator: sense the Zagreb air. 22nd International Conference on Software, Telecommunications and Computer Networks, pp. 423–424.

Antonic, A., Roankovic, K., Marjanovic, M., Pripuic, K., and Zarko, I.P. (2014b) A mobile crowdsensing ecosystem enabled by a cloud-based publish/subscribe middleware. 2014 International Conference on Future Internet of Things and Cloud (FiCloud), IEEE, pp. 107–114.

Banks, A. and Gupta, R. (2014) MQTT Version 3.1.1, OASIS Standard.

Bera, S., Misra, S., and Rodrigues, J.J.P.C. (2015) Cloud computing applications for smart grid: a survey. *IEEE Transactions on Parallel and Distributed Systems*, **26**(5), 1477–1494.

Bonomi, F., Milito, R., Zhu, J., and Addepalli, S. (2012) Fog computing and its role in the Internet of Things. Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing, ACM, pp. 13–16.

Borgia, E. (2014) The Internet of Things vision: key features, applications and open issues. *Computer Communications*, **54**, 1–31.

Brown, B. (2015) Why Micro Datacenters really matter to mobile's future.

Bruneo, D. et al. (2016) Stack4Things as a fog computing platform for Smart City applications. IEEE Conference on Computer Communications Workshops, C IEEE, San Francisco, CA.

Delmastro, F. (2012) Pervasive communications in healthcare. *Computer Communications*, **35**(11), 1284–1295.

Fang, X., Misra, S., Xue, G., and Yang, D. (2012) Managing smart grid information in the cloud: opportunities, model, and applications. *IEEE Network*, **26**(4). doi: 10.1109/MNET.2012.6246750

Garcia Lopez, P. et al. (2015) Edge-centric computing: vision and challenges. *SIGCOMM Computer Communication Review*, **45**(5), 37–42.

Gazis, V., Leonardi, A., Mathioudakis, K., Sasloglou, K., Kikiras, P., and Sudhaakar R. (2015) Components of fog computing in an industrial Internet of Things context. 12th Annual IEEE International Conference on Sensing, Communication, and Networking: Workshops (SECON Workshops), IEEE.

Gia, T.N., Jiang, M., Rahmani, A.-M., Westerlund, T., Liljeberg, P., and Tenhunen, H. (2015) Fog computing in healthcare Internet of Things: a case study on ECG feature extraction. IEEE International Conference on Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM), IEEE.

Gubbi, J., Buyya, R., Marusic, S., and Palaniswami, M. (2013) Internet of Things (IoT): a vision, architectural elements, and future directions. *Future Generation Computer Systems*, **7**(29), 1645–1660.

Happ, D. and Wolisz, A. (2017) Towards gateway to cloud offloading in IoT publish/subscribe systems. 2nd International Conference on Fog & Mobile Edge Computing, IEEE, Valencia, Spain.

Happ, D., Karowski, N., Thomas, H., Menzel, V., and Wolisz, A. (2015) Meeting IoT platform requirements with open pub/sub solutions. 1st International Conference on Cloudification of the Internet of Things (CIoT'15), Paris, France.

Hassan, Q.F., Riad, A.M., and Hassan, A.E. (2012) Understanding cloud computing. *Software Reuse in the Emerging Cloud Computing Era*, IGI Global, pp. 204–227.

Hossain, M.S., Muhammad, G., Alhamid, M.F., Song, B., and Al-Mutib, K. (2016) Audio-visual emotion recognition using big data towards 5G. *Mobile Networks and Applications*, **21**(5), 753–763.

Hunkeler, U., Truong, H.L., and Stanford-Clark, A. (2008) MQTT-S: a publish/subscribe protocol for Wireless Sensor Networks. 3rd International Conference on Communication Systems Software and Middleware and Workshops (COMSWARE'08), Bangalore, India, pp. 791–798.

Jayaraman, P.P., Gomes, J.B., Nguyen, H.L., Abdallah, Z.S., Krishnaswamy, S., and Zaslavsky, A. et al. (2014) East European Conference on Advances in Databases and Information Systems, Springer, Cham.

Kan, C., Chen, Y., Leonelli, F., and Yang, H. (2015) Mobile sensing and network analytics for realizing smart automated systems towards health Internet of Things. IEEE International Conference on Automation Science and Engineering (CASE), IEEE, pp. 1072–1077.

Mahmood, A., Amin, M.B., Hussain, S., Ho Kang, B., Cheong, T., and Lee, S. (2016) Health fog: a novel framework for health and wellness applications. *The Journal of Supercomputing*, **72**(10), 3677–3695.

Mehmood, R.M. and Lee, H.J. (2016) A novel feature extraction method based on late positive potential for emotion recognition in human brain signal patterns. *Computers & Electrical Engineering*, **53**, 444–457.

Mell, P. and Grance, T. (2011) The NIST Definition of Cloud Computing, Computer Security Division, Information Technology Laboratory, National Institute of Standards and Technology, Gaithersburg.

Menzel, T., Karowski, N., Happ, D., Handziski, V., and Wolisz, A. (2014) Social sensor cloud: an architecture meeting cloud-centric IoT platform requirements. 9th KuVS NGSDP Expert Talk on Next Generation Service Delivery Platforms, Berlin, Germany.

Osanaiye, O., Chen, S., Yan, Z., Lu, R., Choo, K.K.R., and Dlodlo, M. (2017) From cloud to fog computing: a review and a conceptual live VM migration framework. *IEEE Access*, **5**, 8284–8300.

Peddi, V.B., Kuhad, P., Yassine, A., Pouladzadeh, P., Shirmohammadi, S., and Shirehjini, A.A.N. (2017) An intelligent cloud-based data processing broker for mobile e-health multimedia applications. *Future Generation Computer Systems*, **66**, 71–86.

Peralta, G., Iglesias-Urkia, M., Barcelo, M., Gomez, R., Moran, A., and Bilbao, J. (2017) Fog computing based efficient IoT scheme for the Industry 4.0. IEEE International Workshop of Electronics, Control, Measurement, Signals and Their Application to Mechatronics, IEEE.

Rusitschka, S., Eger, K., and Gerdes, C. (2010) Smart grid data cloud: a model for utilizing cloud computing in the smart grid domain. *First IEEE International Conference on Smart Grid Communications*, doi: 10.1109/SMARTGRID.2010.5622089.

Sanchez, L. et al. (2014) SmartSantander: IoT experimentation over a smart city testbed. *Computer Networks*, **61**, 217–238.

Sood, S.K. and Mahajan, I. (2017) Wearable IoT sensor based healthcare system for identifying and controlling chikungunya virus. *Computers in Industry*, **91**, 33–44.

Sun, E., Zhang, X., and Li, Z. (2012) The Internet of Things (IOT) and cloud computing (CC) based tailings dam monitoring and pre-alarm system in mines. *Safety Science* **50**(4), 811–815.

Tang, B., Chen, Z., Hefferman, G., Wei, T., He, H., and Yang, Q. (2015) A hierarchical distributed fog computing architecture for big data analysis in smart cities. Proceedings of the ASE Big Data & SocialInformatics, ACM.

Tao, F., Cheng, Y., Xu, L.D., Zhang, L., and Li, B.H. (2014) CCIoT-CMfg: cloud computing and Internet of Things-based cloud manufacturing service system. *IEEE Transactions on Industrial Informatics*, **10**(2), 1435–1442.

Tschofenig, H., Arkko, J., McPherson, D., Thaler, D., and McPherson, D. (2015) Architectural Considerations in Smart Object Networking.

Yi, S., Hao, Z., Qin, Z., and Li, Q. (2015) Fog computing: platform and applications. Third IEEE Workshop on Hot Topics in Web Systems and Technologies (HotWeb), pp. 73–78.

Zhang, B. et al. (2015) The cloud is not enough: saving IoT from the cloud. 7th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud '15), Santa Clara, CA, USENIX Association.