

A Fog Operating System for User-Oriented IoT Services: Challenges and Research Directions

Nakjung Choi, Daewoo Kim, Sung-Ju Lee, and Yung Yi

Fog computing brings computing, storage, and networking even closer to end users and devices for services with better QoS. We introduce FogOS, a fog computing architecture for IoT services. The authors take the perspective of designing an operating system, practicing the architectural lessons from the long history of operating systems.

ABSTRACT

As the proliferation of mobile devices has ignited cloud computing, it is expected that increasing development and deployment of IoT services will expedite the era of fog computing. Fog computing brings computing, storage, and networking even closer to end users and devices for services with better QoS. We introduce FogOS, a fog computing architecture for IoT services. We take the perspective of designing an operating system, practicing the architectural lessons from the long history of operating systems. We focus on addressing the challenges raised by the diversity and heterogeneity of IoT services and edge devices that are owned by individuals and different owners, and presenting how FogOS is designed to effectively and efficiently provide and manage such IoT services. We provide a city-scale surveillance use case to demonstrate FogOS in action.

INTRODUCTION

The Internet of Things (IoT) is no longer a vision, but a reality. We are already witnessing and experiencing many interesting IoT applications. Gartner predicts that about 21 billion “things” across industry sections will be connected to the network by 2020 [1]. These devices generate and transmit data that have diverse requirements in terms of not only volume, but also variety and velocity.

With the ever increasing number of devices and data generated from the edge, the classical cloud-based computing paradigm is faced with challenges, as IDC estimates that the amount of data analyzed on the IoT that are physically at or near the devices is approaching 40 percent [2]. To address these networking and computing trends, *fog computing* brings the cloud closer to the “things” that produce and act on IoT data as depicted in Fig. 1. It is an architectural paradigm that is more appropriate for the fast-growing IoT as it brings computing, storage, and networking closer and faster to the edge devices. We propose Fog Operating System (FogOS), a fog computing architecture for IoT services. We view the whole IoT ecosystem as a computer, and take the perspective of an operating system for our FogOS architecture. The role of traditional operating systems in computers is managing computer hardware and software resources and providing common services for computer programs. FogOS, on the other hand, regards IoT applications (that

correspond to programs in OS) as X-as-a-service (XaaS, e.g., lighting-as-a-service, temperature-sensing-as-a-service) for which common interfaces are provided. The set of resources managed by FogOS include all fog and cloud (e.g., nano/edge cloud) and edge devices (e.g., sensors/actuators). They can also be connected to any level of cloud (e.g., central/regional/metro) when FogOS has an appropriate peering contract.

We particularly consider the case where edge devices are possibly owned by different individuals and providers. Many current IoT devices are deployed by infrastructure providers, but we argue that many more sensors/actuators will be increasingly individually owned and shared in the future when they are appropriately incentivized. In that scenario, there is a complex economic interplay between different players (e.g., IoT users, IoT application providers, infrastructure providers, and edge device owners). Hence, FogOS can function as a distributed operating system that manages the cloud and the resources at the edge, and a platform of incentivizing and connecting individually owned edge devices.¹

An OS for network resource management, such as Open Network Operating System (ONOS) [3], is a seemingly similar concept to FogOS. For example, ONOS has been proposed as a control platform of software defined networking (SDN) for carrier and cloud provider networks, with scalability, availability, and performance in mind. However, there are two main differences between FogOS and ONOS.

First, ONOS operators directly own and control all network devices that are fixed with relatively stable operating conditions, while FogOS needs to control significantly diverse edge devices that are highly dynamic and owned by different parties. Thus, FogOS needs to play the role of a broker of pooling/slicing edge devices’ resources and coordinating all the players with self-interest. Second, ONOS participates in standardizing device interfacing (e.g., OpenFlow [4]). However, in IoT, diversity in devices, services, and protocols is inevitable, resulting in a more challenging environment. A cloud computing OS, such as OpenStack [5], also aims to orchestrate the large-scale computing resources in a shared infrastructure built on top of standard and commodity hardware under the same administrative domain. Hence, FogOS has similar major differences compared to the traditional cloud computing OS.

There are three groups actively designing

¹ Fog: typically means both clouds at edge and user devices in the literature, but throughout this article, we use the terms “fog cloud” and “edge device.”

architectures for fog computing: the Open Fog Consortium [6], the European Telecommunications Standards Institute's (ETSI's) mobile edge computing (MEC) [7], and cloudlets [8], each with slightly different visions and emphasis (see [9] for comparison). We believe that FogOS can be applied to or even merged with any of these architectures, as our focus is on handling the diversity and heterogeneity of user-oriented IoT services and edge devices that are owned by individuals and different owners using fog computing.

FOG OPERATING SYSTEM: ARCHITECTURE

KEY CHALLENGES

We describe the challenges of fog computing architecture for highly diverse IoT applications with heterogeneous edge devices owned by different individuals and providers (see Table 1 for a summary and existing solutions).

Scalability: Being at exponential growth, there would be a significant number of IoT devices, which in turn run various IoT applications and generate a sheer amount of data.

Complex Inter-Networking: Due to the large scale and diversity, IoT devices will be physically connected in various forms and under diverse conditions, for example, wireless multihop connectivity using heterogeneous radio access technologies, often with mobility.

Dynamics and Adaptation: With wireless connectivity and mobility, IoT devices experience frequent environmental changes in topology and communication conditions. In addition, IoT applications may have diverse lifetimes and quality of service (QoS) requirements, requiring prompt allocation of edge resources and re-embedding of IoT applications.

Diversity and Heterogeneity: Edge devices have various capabilities in communication radios, sensors, computing powers, storage, and so on. This requires seamless interfacing and interoperability, often incurring non-negligible overhead and yielding implementation/operation complexity.

In FogOS, we tackle the above challenges using a reference architecture, as depicted in Fig. 2, consisting of the following four main components:

- Service and device abstraction
- Resource management
- Application management
- Edge resource: registration, ID/addressing, and control interface

The challenges due to diversity and heterogeneity are resolved by an abstraction layer for services and devices (see the following subsection). The application and resource managers work closely together to provide complex internetworking services and adaptively allocate edge/fog resources to accommodate the dynamics of applications and resources (see "Resource Management" and "Application Management" below). In the "Edge Resource: Registration, Identification, and Control Interface" section, we describe ways of improving network and service scalability.

SERVICE AND DEVICE ABSTRACTION

In the fog computing environment where FogOS operates, there is a common property in IoT applications and edge devices: diversity. This diversi-

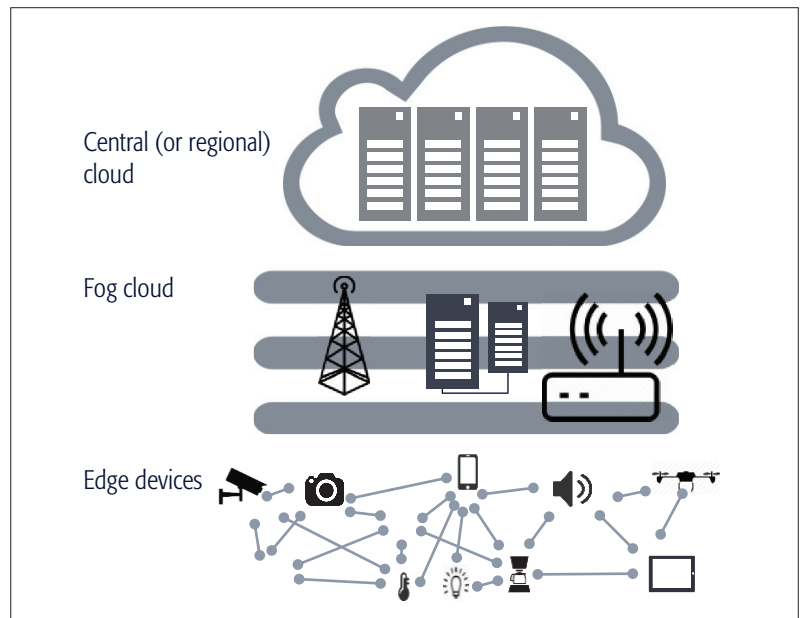


Figure 1. Fog cloud and edge devices.

ty complicates the process of developing an IoT application and the control of edge devices. It is therefore necessary to provide flexible but consistent abstraction as application programming interfaces (APIs), both from FogOS to applications (service abstraction or service API) and from FogOS to edge devices (device API). These APIs are designed and categorized by the degree of generalization and exposure, where generalization refers to how concrete abstraction should be, and exposure deals with how controllable we should make service and device through the designed APIs.

Service Abstraction

Generalization: In operating systems, users can directly access OS resources by invoking low-level system calls or using a high-level programming language dependent standard library. FogOS defines the following three hierarchical service APIs from low to high level:

- *Level 1: Resource service API:* This API resides at the lowest level, providing services that can control each individual edge device resource such as computing, storage, sensing, actuating, and radio access (or link) to applications. For example, a service call of "read the temperature from sensor X" can be invoked by an application.
- *Level 2: Network service API:* Using a collection of resource service APIs, this API provides the service of creating a networked slice that consists of some set of edge device resources. For example, to create a video surveillance service, a service call of "form a wireless video sensor network with video sensors X, Y, Z and an edge cloud C, where all wireless sensor nodes are connected to C with statistical bandwidth guarantee of 1 Mb/s." Note that there can be a multihop path from X to C, where the resource matching module of the service manager determines the "optimal" path (see "Application Management" later in the article).

Components (FogOS)	Task	Example func. in general OS	Challenges	Existing solutions
Service and device abstraction	Providing a device data model	File as a universal resource identification	Diversity in devices	Device data model in [10, 11]
	Providing service APIs	System calls and standard libraries	Diversity in services	IoT service APIs (e.g., IoBridge, Evrythng)
Resource management	Pooling resources	Distributed system (e.g., Hadoop distributed file system)	Spatially separated resources, heterogeneous devices	Network controller for SDN (e.g., ONOS [3])
	Slicing resources	Virtual memory	Lightweight slicing	Hypervisor for compute resource
Application management	Matching services with resources	CPU scheduling	Adaptation to dynamics environments (e.g., diverse service lifetime, devices' mobility)	Virtual network embedding and adaptation
Edge resource control	Registration and identification	Device manager	Diversity and large scale	AllJoyn and IoTivity's administration system [10, 11]
	Control interface	System bus	Heterogeneous network interfaces	Openflow, CoAP, MQTT

Table 1. Challenges and solutions: fog computing.

- **Level 3: Application service API:** This API is at the highest abstraction and allows application developers to easily create a service that is defined as a typical service a priori. An example could be a call of “create a video surveillance service at hotspots with high-definition TV quality.”

Note that there could be more levels in this hierarchy. Application developers are allowed to utilize any level of service APIs with different controllability and programming proficiency.

Exposure: A FogOS designer may choose different exposure degrees even in each service API, based on programming friendliness and security level. For example, a video surveillance service can be created with the service requirement description {3 cameras} or {1 camera near *gps-1*, 2 cameras near *gps-2*} at the level 3 resource API. Also, everything cannot be open to application developers. For example, network link resources in the resource API at level 1 cannot be accessible because arbitrary change of link resources may negatively affect other applications; even the entire resource API at level 1 can be blocked to allow only high-level access.

Device Abstraction

Generalization: UNIX-like operating systems treat everything as a file, for example, */dev/sda1* for a hard disk, */tmp/mongodb-27017.sock* for a socket, and */proc* filesystem (*procfs*) for a process or other system information. FogOS could enjoy a similar level of generalization, but has more diversity to interface with various existing and emerging edge devices, which are possibly manufactured by different vendors. To this end, multiple device data models are defined and exposed to FogOS, for example, “sensor device data model,” which is different from a single device data model in operating systems (i.e., a file). Note that device data models might have inheritance relation as in the object oriented programming language; for example, a temperature sensor device data model inherits a sensor device data model, or some device data models

might be grouped. This generalized abstraction of edge devices enables emerging IoT devices to easily be incorporated without affecting existing applications.

Exposure: Typical OSs provide different control granularity to each managed resource. For example, a default WiFi device driver provides the interface to configure WiFi behaviors. However, a vendor-provided device driver can be activated to expose vendor-specific features with finer control granularity. Similarly, edge devices that are in the same device category might have their own specific features that can be differentiated from other edge devices. Hence, FogOS still requires vendor-specific/owner-specific device drivers to control devices' details or new features, in addition to general and abstract IoT device data model-based control.

RESOURCE MANAGEMENT

FogOS manages the resources of edge devices and fog clouds that are spatially separated and often need to be controlled in a distributed manner. We assume that the list of available resources are registered at the resource management module for the process of edge resource registration. As in traditional OS, FogOS pools or slices the available resources whenever needed, but there are many challenges to be handled, as elaborated next.

Resource Pooling: The concept of resource pooling is used in a variety of contexts across different domains. In this article, we define resource pooling as a mechanism to collect the resources of the same “class.” A good example in a general OS is the notion of virtual memory in the hierarchical memory system, which enables the main memory and hard disk to be pooled, transparently seen as runtime storage by running processes.

In fog computing, similar resource pooling would be useful in furnishing IoT applications with larger service options and freedom. The unique challenges of resource pooling in fog computing are:

- Pooling occurs among edge devices that might be placed in spatially different locations.




Figure 2. FogOS (Fog Operating System) reference model.

- The resources to be pooled are highly heterogeneous.
 - Limited resources of edge devices often require large-scale pooling.
- These unique features act as technical challenges that should be tackled by FogOS. A candidate list of resource pooling is as follows.

Computing/Storage Pooling: The processing power of an edge device is likely limited. For intense data processing that requires fast response, we can use multiple edge devices in a distributed manner. Similarly, limited storage can be compensated by a distributed collection of storage of other devices.

Sensor/Actuator Pooling: Many IoT applications relying on data from sensors might increase information accuracy by exploiting similar data from multiple sensors. Also, different kinds of sensors lead to more complete information on the monitoring status, for example, pooling and fusion of the data from a combination of gyroscopes, magnetometers, and accelerometers. In actuator pooling, a good example is multiple drones flying

in a group, performing environmental sensing in a collaborative manner.

Network Link Pooling: IoT applications that generate large data or require low latency need high-speed access links. To that end, an edge device with multiple communication radios can pool them to create a thick communication pipe. Also, a device that is not directly connected to a fog cloud could use other edge nodes as relays.

Resource Slicing: As opposed to resource pooling, resource slicing corresponds to a mechanism that enables sharing of physical resources by multiple IoT applications. For example, in an OS, storage can be sliced through the concept of virtual memory space so that multiple processes can regard the entire (virtual) memory as if it is exclusively allocated to each single process. Multicore CPU allocating each core independently to each process is another example.

Slicing of the resources of edge devices can provide differential granularity and help use the resources efficiently. Sensor/actuator and processing resources can be sliced temporally, and stor-

We assume that the list of available resources are registered at the resource management module for the process of edge resource registration. As in traditional OS, FogOS pools or slices the available resources, whenever needed, but there exist many challenges to be handled, as elaborated next.

Different IoT applications need different types and amount of edge resources, depending on their QoS requirements. One of key functions of the application manager is to compute a matching solution from such requirements to the edge resources, where the available resources are obtained by querying to the resource manager.

age resource can be sliced spatially. The network link resource can also be sliced temporally as well as spatially (e.g., subdivision of the communication channel). Synchronization and scheduling among the distributed resources is a key challenge, as poor execution would result in serious resource waste.

APPLICATION MANAGEMENT

As the resource manager manages all the edge device resources, the application manager manages everything on the running IoT applications by matching the service requests to the edge resources, monitoring the running application's resource usage status and enforcing service level agreements (SLAs), orchestrating the registered and available edge resources among multiple ones (e.g., prioritization), concurrent applications, and adapting to the changes of edge resource and application status.

Application-Edge Resource Matching: Different IoT applications need different types and amounts of edge resources, depending on their QoS requirements. One of the key functions of the application manager is to compute a matching solution from such requirements to the edge resources, where the available resources are obtained by querying the resource manager. In a typical OS, such a matching is trivial as the device resources are directly controllable and small scale. However, in IoT, there are many and diverse devices, often placed in spatially different locations, requiring networked control. Diversity requires the matching module to match the required resource "optimally" from many candidates. Depending on how effective this matching is, the number of IoT applications that can be accepted and run is determined, which has large impact on the revenue of IoT service providers and other economic players. Theoretical understanding of this matching problem must be made, and practically implementable algorithms with low complexity are of significant importance, which in turn depends on the type of applications provided as IoT application service APIs.

Policing, Scheduling, and Orchestration: Once edge resources are appropriately allocated to incoming IoT applications, the application manager keeps track of their resource usage and monitors whether SLAs are violated. SLA violation of an application might degrade QoS of other applications, for which a certain level of resource partitioning often becomes of some value. The key challenge comes from the large scale of edge devices, where monitoring and policing for each would incur significant overhead even for a small edge device.

Dynamic creation and termination of IoT applications fluctuates available edge resources over time, often leading to the resource competition among them. In addition, each application would be assigned a different priority based on, for instance, security and pricing. All these motivate FogOS to employ a smart scheduler of running applications, similar to the job scheduler of an OS.

Adaptation to Resource and Service Changes: After the applications are matched to a set of edge resources, this matching result might not remain valid due to the change of application status and

the change/fault of edge resources. Whenever applications with higher priority arrive or existing applications terminate, the application manager might need to re-match the resource among the running applications for better resource management. This adaptation is also necessary when the edge resource topology changes as edge devices move or experience fault (e.g., battery shortage). In this case, the application manager must support continuous reconfiguration of the application and edge resources in collaboration with the resource manager. However, we must consider the trade-off between operation cost efficiency of reconfiguration and performance.

EDGE RESOURCE: REGISTRATION, IDENTIFICATION, AND CONTROL INTERFACE

Identification and Addressing: In fog computing, high dynamics and diversity of edge networks force FogOS to interact with edge resources frequently to keep an up-to-date snapshot of the resource store of the resource manager (Fig. 2), and pool/negotiate a proper set of resources to embed various IoT applications, where an efficient identification of edge resources via IDing and addressing is essential. We propose to use both syntactic and semantic IDs in FogOS. Syntactic IDs refer to the ones that directly identify the edge resource (e.g., a 5th sensor of room 2 of building 2 of the Korea Advanced Institute of Science and Technology, KAIST), whereas semantic IDs support the context of what a service want to utilize (e.g., any temperature sensor sensing 10°C of room 2 of building 2). Each ID might use a different binding with its network address, for example, static binding for syntactic IDs and dynamic binding for semantic IDs.

One can refer to IoTivity's identification specification [10], AllJoyn [11], geocasting [12], or Named Data Networking (NDN) for IoT [13]. For example, AllJoyn requires each IoT device to know the minimum information (i.e., name) of destinations, and each device can find the destination device with this information by using mDNS (multicast). However, while AllJoyn is suited for a small-scale IoT network, FogOS targets networks with a large number of edge devices with high dynamics and diversity, and hence requires scalable solutions.

Resource Discovery, Registration, and Management: FogOS must discover edge devices and their resources, manage the list, and monitor their status. Two schemes are possible: proactive and reactive. In a proactive scheme, when an edge device enters our FogOS-administered network, it notifies FogOS of its intention to join with its list of available resources. FogOS then updates its resource store database to keep track of this new edge resource. In order to keep the information up to date, the available resource status must be periodically reported to the resource manager of FogOS. In a reactive scheme, edge resources are queried on demand, whenever new edge resources are needed as new applications are about to be created. Proactive schemes provide faster response to resource lookup and matching for new applications, but at the cost of larger overhead stemming from keeping track of the resource-related information. On the other hand, reactive schemes provide fresher information but

with slower response time. AllJoyn follows this reactive scheme, mainly because it is designed for home-scale one-hop IoT applications. We envision FogOS operating on a larger scale; thus, we believe that a scheme with a certain degree of proactivity is necessary for a possible hybrid approach.

Heterogeneous Control and Network Protocols: As discussed earlier, FogOS uses device APIs to control each individual edge device and fog cloud. We argue that to control fog clouds, the current approach to SDN (i.e., OpenFlow) is a good option. However, to control diverse resource-constrained edge devices, the classical SDN approach might be too heavy and inflexible. Thus, a lightweight version of SDN could be a candidate solution for separating data and control planes. Many of the challenges are due to high heterogeneity in control, communication, and networking protocols of edge devices. The control plane should leverage existing IoT control protocols such as Constrained Application Protocol (CoAP) and Message Queuing Telemetry Transport (MQTT), and also emerging architectures such as information-centric networking (ICN). Similarly, the data plane should support diverse wireless technologies, for example, WiFi, LTE, Low Power WAN (LPWAN) [14], and ZigBee, which is necessary to deliver data as well as control information. There are different proposals for this, where one is to employ a gateway that can understand such heterogeneity, but such an approach of only a single hop at the last mile might restrict the service coverage, and thus limiting the scope of possible IoT applications that FogOS supports. To extend the reach of FogOS, seamless multihop communication over a large-scale wireless nodes would be highly valuable.

FOGOS-DRIVEN IOT ECOSYSTEM

In the fog computing market, there are four key economic players that compete and cooperate to increase their revenues. This is depicted in Fig. 3. We do not claim that the ecosystem mentioned in this article is the only one that would emerge. Rather, we believe that it might be one of the most basic and intuitive patterns where players interact.

End Service² Users (SUs): These are end users who are ready to enjoy IoT applications. They pay the application service fee to service providers under a variety of tariffs.

Edge Resource Owners (EROs): These are individuals or large companies (e.g., mobile network operators that have large-scale communication and sensor infrastructures) who own edge resources or fog clouds. In particular, individual edge resource owners share their resources and partially or entirely sell the resources to an infrastructure provider (InP). They act similarly to Uber drivers. They need to be appropriately incentivized to share the resources, where the incentive mechanism would be given by InPs.

Service Providers (SPs): SPs create diverse IoT applications that attract SUs as over-the-top (OTT) providers. Logically, they do not own the resources of fog clouds or edge devices, but rent them. Thus, they make a contract with InPs that manage the edge resources. Note that it is possible that SPs and InPs are run by a single company. Appli-




Figure 3. Major economic players in fog computing and their interaction patterns.

cation development is made based on the service APIs opened by FogOS.

Infrastructure Providers: They are the ones who run FogOS. InPs have infrastructure of edge devices and fog clouds, and might rely on individual EROs for a large portion of IoT infrastructure. They are required to develop a nice incentive mechanism to attract as many EROs as possible at low cost. Their resources interface with FogOS via device APIs, and they sell their owned and leased resources to SPs. They make profit through business with SPs and EROs.

Note that this market is open to many diverse competition and cooperation scenarios. Edge device owners may act selfishly to maximize their individual revenue, or cooperate with InPs under fair revenue sharing mechanisms. As mentioned earlier, some big player might behave as multiple players. For example, an InP such as a mobile network operator that already has a large-scale cellular and WiFi infrastructure minimally relies on edge device owners by deploying city-scale or even nation-wide sensor/actuator platforms; in such cases, big players are highly likely to provide IoT applications as well to make a large profit from the IoT industry. Non-cooperative and cooperative game theory helps understand the complex interplay in this market and predict the business landscape.

USE CASE AND DEMONSTRATOR: CITY-SCALE SURVEILLANCE SERVICE

We now present the use case of FogOS, a large-scale surveillance service, where Fig. 4 shows our preliminary proof-of-concept implementation of FogOS for a drone-based surveillance service.

SCENARIO OVERVIEW

This is an example of sensing-as-a-service that deploys city-scale surveillance, originally starting with a set of sensing of some target regions, and extends to the service coverage change with drone-driven moving sensors. In this case, an SP can be an IoT sensor service provider, and SUs may include a public safety agency.

Original Service:

- An SP requests a surveillance service to an InP (running FogOS) with a service requirement description $\{K$ regions, M videos, N audios, P sensors $\}$ through a level 3 applica-

Proactive schemes provide faster response to resource look-up and matching for new applications, but at the cost of larger overhead stemming from keeping track of the resource-related information. On the other hand, reactive schemes provide fresher information but with slower response time.

² In earlier sections, we used the term “IoT applications” rather than “IoT service” as “service” is also used as the service provided by FogOS to applications. However, in this section, we use service to mean IoT application service, unless otherwise noted.




Figure 4. Example implementation of surveillance service on FogOS: extension of drone-based moving video sensing.

- tion service API.
- FogOS searches available edge resources owned by itself as well as those owned by EROs (possibly with different priorities), and let the application manager allocate the required resources to embed this application (i.e., matching). These resources might be ready in advance through a proactive resource discovery mechanism, or be searched through a reactive mechanism as discussed earlier.
- The application manager's matching algorithm produces an embedding solution, allocates the computed resource by communicating with the resource manager to retrieve existing/available fog/edge resource information, and commands the resource manager to perform the allocation action.
- SUs enjoy this surveillance service.
- The resource manager periodically monitors the resource usage at associated edge devices, and collaborates with the service manager whenever there is any change or fault of the existing edge resources.

SERVICE EXTENSION

- The SP intends to observe more details of a specific region, say R due to an expected crime, for instance. It requests to add a drone-based video sensing of region R with the modified service description (region R , 1 video sensing with drone, AVAILABLE sensors through level 1 and 2 resource and network service APIs. This service corresponds to live video streaming at the edge cloud in region R to a single or multiple SUs.
- FogOS searches its resource pool, and finds a fog cloud as well as a group of WiFi APs by the InP, but failed to find a drone. It broadcasts a request to find a drone with a video sensor to the EROs in R .
- The video scenes captured by drones and sensing data from the original service reach the allocated fog cloud in region R , which performs augmented reality (AR) functions to generate a richer content of the scene view. This post-processed video stream is delivered to multiple SUs.

PROOF-OF-CONCEPT IMPLEMENTATION OF FOGOS

To see the feasibility of FogOS, we implement a prototype, where FogOS plays the following two roles: controller and platform for IoT ecosystem. In our implementation, the economic interaction between key players is simplified, that is, when EROs register their resources to FogOS, EROs' resources are shared through InP. In this article, we mainly focus on the control function of FogOS, as follows:

- Drones and sensors are controlled by an application running on the FogOS through service and device abstraction layers. Thus, we are able to control flying drones and SDN IoT sensors through FogOS.
- Computing, sensing, and networking resources are pooled together and matched to this service by the resource and service managers of FogOS.
- A video from drones and sensing data is processed/merged by the allocated computing resources, and then multiple views for different SUs are created, as shown in Fig. 4.

CONCLUSION

We introduce a fog computing and networking architecture for IoT services, termed FogOS, practicing architectural lessons from operating systems. FogOS is composed of four major components: service/resource abstraction, resource manager, application manager, and edge resource identification/registration, whose challenges and main research directions are discussed. We hope that our vision in FogOS will be shared by other groups in academia and industry working on IoT and fog computing, and more constructive discussions will continue to follow, inspired by FogOS. These future directions include the extension of FogOS to support the key scenarios in the fifth generation, that is, enhanced mobile broadband, ultra-reliable and low-latency communications, and massive machine type communications.

ACKNOWLEDGMENT

This work was partially supported by the Institute for Information and Communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No.B0717-17-0034, Versatile Network System Architecture for Multi-Dimensional Diversity).

REFERENCES

- [1] Gartner, "Gartner Says 6.4 Billion Connected Things Will Be in Use in 2016, Up 30 Percent from 2015"; <http://www.gartner.com/newsroom/id/3165317>, accessed 4 May 2017.
- [2] V. Turner *et al.*, "IDC FutureScape: Worldwide Internet of Things 2015 Predictions," Int'l. Data Corp., 2014.
- [3] ON.LAB. "ONOS — A New Carrier-Grade SDN Network Operating System Designed for High Availability, Performance, Scale-out"; <http://onosproject.org>, accessed 4 May 2017.
- [4] ONF, OpenFlow Switch Specificatio, v. 1.5.1, 2015.
- [5] "OpenStack: Open Source Software for Creating Private and Public Clouds"; <https://www.openstack.org/>, accessed 4 May 2017.
- [6] Open Fog Consortium; <https://www.openfogconsortium.org/>, accessed 4 May 2017.
- [7] M. Patel *et al.*, "Mobile-Edge Computing Introductory Technical White Paper," MEC Industry Initiative, 2014.
- [8] M. Satyanarayanan *et al.*, "Cloudlets: At the Leading Edge Of Mobile-Cloud Convergence," *Proc. IEEE Mobile Computing, Applications and Services*, 2014.
- [9] G. I. Klas, "Fog Computing and Mobile Edge Cloud Gain Momentum Open Fog Consortium, ETSI MEC and Cloudlets," 2015.
- [10] OCF, "OIC Core Candidate Specification," 2016.
- [11] A. Alliance, "AllJoyn Framework," 2016; <https://allseenalliance.org/framework>, accessed 4 May 2017.
- [12] Y.-B. Ko and N. H. Vaidya, "Geotora: A Protocol for Geocasting in Mobile Ad Hoc Networks," *Proc. IEEE Int'l. Conf. Network Protocols*, 2000.
- [13] E. Baccelli *et al.*, "Information Centric Networking in the IoT: Experiments with NDN in the Wild," *Proc. ACM Information-Centric Networking*, 2014.
- [14] Nokia Networks, "LTE-M - Optimizing LTE for the Internet of Things," white paper, Aug. 2015.

BIOGRAPHIES

NAKJUNG CHOI (nakjung.choi@nokia-bell-labs.com) is a member of technical staff at Nokia Bell Labs, Murray Hill, New Jersey, since April 2010. He received his B.S. (magna cum laude) and Ph.D. at the School of Computer Science and Engineering, Seoul National University in 2002 and 2009, respectively. Also, he has received several awards such as Best Paper Awards and Awards of Excellence. His research is focused on SDN/NFV/cloud, 4G/5G/IoT, and future converged services.

DAEWOO KIM (daewookim@kaist.ac.kr) received his B.S. from the Department of Electrical Engineering, Yonsei University, South Korea, in 2013. He is a doctoral student at the School of Electrical Engineering, KAIST, since 2013. His research interests include sensor networks, network economics, fog computing, and machine learning in networking.

SUNG-JU LEE (profsj@kaist.ac.kr) is an associate professor and KAIST Endowed Chair Professor at KAIST. He received his Ph.D. in computer science from the University of California, Los Angeles in 2000, and spent 15 years in the industry in Silicon Valley before joining KAIST. His research interests include computer networks, mobile computing, network security, and HCI. He is the winner of the HP CEO Innovation Award, the Best Paper Award at IEEE ICDCS 2016, and the Test-of-Time Paper Award at ACM WINTECH 2016.

YUNG YI (yiyung@kaist.ac.kr) is an associate professor at the Department of Electrical Engineering at KAIST. He received his Ph.D. from the Department of Electrical and Computer Engineering, University of Texas at Austin in 2006. His research interests include computer networks and machine learning. He received the best paper awards at IEEE SECON and ACM Mobihoc in 2013, and he was the winner of the IEEE William R. Bennet Prize in 2016.