IoT in the Fog: A Roadmap for Data-Centric IoT Development

Sharief M. A. Oteafy and Hossam S. Hassanein

Abstract

Our interactions with the world are increasingly dependent on context-aware services, and the future of smart cities is coupled with how efficiently and reliably we can deliver these services to end users. In this article we present the premise of personalized IoT systems, by leveraging novel advancements in user-centric technologies under the fog computing architecture. This means leveraging the connectivity and processing potential of the fog to bring IoT control and analytics closer to the user, and improve the coupling of services with local IoT components in user-centric contexts. The potential gain in access latency and context-sensitive service matching will enable a multitude of smart city services. On one hand, data management (collection, pruning, *denaturing* [1], and encryption) can take place closer to the edge, thereby leveraging network load and service times. On the other hand, service matching in smart city applications will witness higher responsiveness and resource visibility in areas with intermittent connectivity or high mobility. We first present the challenges in migrating cloud-IoT architectures to the network edge, and detail the hindrances in transitioning the control and management of IoT systems to the user end. As a remedy, we survey recent advancements in the IoT, ubiquitous computing, and user-centric services, which enable us to advance personalized IoT architectures. We finally present a framework for IoT in the fog to synergize these advancements, and present a proof-of-concept use case to highlight its utility and impact. We conclude this article with prime directions for future work to realize a personalized IoT architecture, and highlight the potential gain in prioritizing five high-yield potential research issues.

Understanding Fog IoT

The case for cloud computing (CC) infrastructures is widely established. Simply put, the ability to offload computationally intensive tasks on remote data centers, where you are *elastically* charged for what you use, is growing as a preferable alternative in a large spectrum of applications. The most prominent everyday use of such services is witnessed in speech recognition software (e.g., Samsung S-voice, Google Talk and Apple's Siri), near-real-time pattern recognition for object identification (e.g., YOLO — You Only Look Once: Unified, Real-Time Object Detection),[1] and translation (e.g., Google Translate).

However, as we transition into a mobile-driven world, today's users are expecting crisp interaction with their surrounding technologies. The user can no longer afford to wait for the typically varying response time of a cloud-based computation or service discovery, especially under intermittent connectivity, high-mobility scenarios, or with stringent demands on tolerated delay. The rising tide of improving quality of experience (QoE) as well as enabling contextualized user-centric applications is driving forward the migration toward fog computing.

This user-driven shift in computation and storage to near-edge fog architectures is enabling many applications that require less interaction with remote services (or data centers). Fog computing builds on research in edge analytics [2] and leverages recent developments in cloudlets [1]. The interplay between the cloud, cloudlet/edge, fog, mist, and end users is depicted in Fig. 1.

More importantly, as the Internet of Things (IoT) bridges the physical and virtual worlds of interactions, we need solutions that contextualize our interactions with immediate resources. That is, we now have the technology to probe surrounding resources (sensing, processing, communication, etc.) in real time [3], but lack the framework to deliver a responsive and user-centric IoT experience on the go.

In this article we survey the challenges in realizing IoT systems in the fog, and present an overview of recent advancements that could be synergized to deliver a personalized IoT ecosystem. We target a framework that will encompass heterogeneous IoT resources in a given region, and the variation in processing and communication offloading that could be leveraged by the fog IoT architecture.

Fog Computing

In its simplest terms, the fog layer resides between a local resource and the cloud service. In theory, a multiplicity of fog nodes will be geographically distributed to service local resources in their respective regions. This multi-tiered architecture is depicted in Fig. 1. Each of these nodes will be able to leverage computing tasks and take part in service orchestration with underlying resources in the region [4].

The authors first present the challenges in migrating cloud-IoT architectures to the network edge, and detail the hindrances in transitioning the control and management of IoT systems to the user-end. As a remedy, they survey recent advancements in the IoT, ubiquitous computing, and user-centric services, which enable them to advance personalized IoT architectures.

[1] https://pjreddie.com/publications/yolo

Sharief M. A. Oteafy is with DePaul University; Hossam S. Hassanein is with Queen's University.

Digital Object Identifier:
10.1109/MCOM.2018.1700299

0163-6804/18/$25.00 © 2018 IEEE

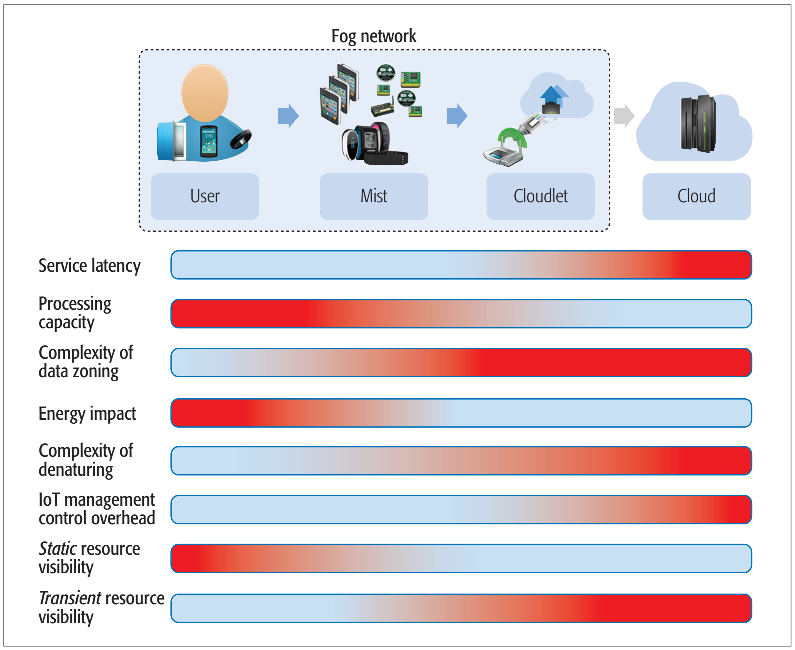
Fog network
User
Mist
Cloudlet
Cloud
Service latency
Processing capacity
Complexity of data zoning
Energy impact
Complexity of denaturing
IoT management control overhead
Static resource visibility
Transient resource visibility

Figure 1. Overview of tiers in a cloud-IoT architecture, highlighting the span of fog networks. The major challenges in realizing IoT operation are depicted in bars below each of the architectural components, highlighting the variation from simple/better (in light blue) to complex/worse (in red) for each of the operational mandates/design challenges. In some scenarios, as in energy impact on IoT resources, the variation is non-increasing in either directions, but exhibits better results under a subset of the tiers (cloudlet/edge tier in this case) [1].

The notion of fog computing is preceded by earlier work on *cloudlet* access, whereby an intermediate connection/access point is deployed to bridge the computational offloading process from mobile devices to cloud services. Earlier work by M. Satyanarayanan, overviewed in [1], presents a detailed account on the motivation behind cloudlet design, and highlights its soft-state that is inherently more flexible, modular, and distributed in contrast to cloud platforms.

Variants of Cloud-Based Sensing

The notion of leveraging cloud sensing has been investigated heavily in the past decade [2, 4–6], mainly to enable public sensing schemes. Cloud sensing is mainly concerned with distributed data collection for offline querying [5], and newer models attempt to leverage cloud services to enable a real-time association between service requesters (i.e., application requiring specific data in real-time) and currently available resources that are connected to the cloud sensing architecture [2]. However, most cloud sensing architectures tolerate a significant delay in processing, and are intrinsically designed for offline operation, both of which hinder its application in newer systems where devices are mobile, intermittently reachable, and more invested in real-time information services. In recent developments, the case for mobile edge computing (MEC) and novel technologies that bring more processing to the edge of the network are enabling newer forms of cloud sensing in the nearer fog. This means building systems that manage, disseminate, and respond to queries using in-field technologies rather than *foraging* resources from distant cloud services. This notion of *fog sensing* is at the heart of what this article covers.

Fog-Enabled Services

The projected proliferation of machine-to-machine (M2M) services, along with an evident transition into mobile-driven services and applications, are rendering many cloud-dependent services inefficient and restricting. In addition, the inherent heterogeneity of all devices that are joining the mobile resource pool is increasing the complexity and delay in centralized (cloud-based) management and orchestration of information services over these devices. The advent of big sensed data, in terms of data produced and potential services enabled by the aggregation of all these resources, is further established in [7], and we are in dire need of an architecture that can access, monitor, manage, and recruit these services in real time and within their respective contexts [1]. We next survey the major challenges facing our development of fog sensing, and then present a roadmap to its development in light of novel technologies.

Migrating IoT Services to the Edge: Core Challenges

The sheer amount of IoT/M2M traffic projected in the next five years is mandating novel design considerations in both data processing and communication management. Typical traffic generation in IoT devices in 2016 averaged 1614 MB/month,[1] mostly from wearable devices. However, with a projection[2] of a rise in number of M2M connections from 1.1 billion (in 2017) to 3.3 billion (in 2021), there are many scalability challenges to address.

Earlier research on *cyber foraging* by M. Satyanarayanan argued that regardless of hardware advances at the user end, static resources on the Internet (or distributed systems in the general sense) will remain far superior [8]. Thus, cyber foraging was based on a growing disparity between resource capability at the edge in contrast to that in *cyberspace*. The argument for Internet-based resource foraging grew significantly with the realization of cloud architectures, and most IoT developments attempted to capitalize on resource abundance and elastic pricing of cloud services. However, in revisiting the recent explosion in data usage and stringency of time limits, Satyanarayanan and others have argued for reducing the dependence on "remote" cyber foraging in the cloud.

Migration of data and communication control to the edge of the network has been at the heart of context-aware services for over a decade [5]. Many attempts at *personalizing* IoT interactions have brought control to the edge, mostly at the user device or gateway levels. The benefits in response latency, hub-free M2M interactions, and power conservation have been major drivers of near-edge operation. This approach opened the door for IoT systems that probe nearby resources for real-time service matching [9] and enabling context-aware IoT [6].

While these drivers are indeed pressing, there are many challenges as we migrate

[2] Cisco Visual Networking Index (VNI): Global Mobile Data Traffic Forecast Update, 2016–2021, March 2017 update.

IoT layer Factor	User	Mist	Cloudlet/edge	Cloud
Architectural components	Personal smart devices (phones, tablets, vehicles, wearables, etc.)	Home automation devices (e.g., NEST thermostat), neighboring smart devices, mostly standalone IoT devices	High-end access points with processing, connectivity and storage resources	Ultra-large-scale processing and storage resources, backbone access to Tier-1 Internet
Geo-distribution	Local to user, directly accessible	Non-uniform, often agnostic to cloudlet locations	Close to backbone connections, higher urban density	Demand-centric
Data pruning techniques	Local homogeneous fusion (temporal sampling, averaging, simple flat fusion)	Neighborhood-based data alignment and fusion	Geo-sensitive Hierarchical data fusion and cleaning	Large-scale data classification, mining, and cataloging
Time responsiveness (processing)	Immediate, zero network delay	Typically 1-hop delay, limited queuing delay	Short contention depending on users, mobility, backend buffering, and other factors	Longest delay, depending on congestion toward remote cloud service and aggregated queuing delay
Mobility	Highly mobile	Typically mobile (often as a group, as in VANET)	Typically static, with recent work on vehicle-mounted cloudlets	Strictly static
Deployment plan	Ad hoc	Mostly ad hoc and movable. Static exceptions exist (e.g., home automation)	Mostly coupled with APs and high-BW backbone connections	Strategically placed by CDN and cloud providers (both surrogate servers and data centers)
Networking technologies	Typically "multi-home" to short- and long-range networks	M2M and short-range communication	LTE/5G backbone Ethernet (wired)	Mostly > 10 GbE
Current standards	IEEE 802.11/15 family LTE/4G BLE/Bluetooth NFC EC-GSM-IoT DASH7	NFC BL ZigBee DSRC (vehicular) RPMA	IEEE 802.11ac MulteFire LTE / 5G	Dense wavelength-division multiplexing (DWDM) EN 50600-2-4 TIA-942-A
Governing architectures	Personal applications	M2M H2M	VMware Dedicated IoT hubs	Large-scale data centers (e.g., EC2)
Architectural advantage	Zero delay Immune to mobility challenges Least challenge with security and authentication	Harnessing resources from immediate neighborhood (sensing, communication, processing, etc) Immune to cloud service disruptions Inherently geo-contextualized	Reduces ingress traffic Masks cloud unavailability (due to connectivity, DDoS attacks, etc.) Anonymization techniques	Most economic use of resources Enables large-scale view of resources Most suited for large-scale data analytics

Figure 2. Contrasting the design factors in delivering IoT operation over Cloud variants. More importantly, the multi-tiered approach to Cloud interaction is highlighted over the different levels of user involvement with personal devices, to neighboring IoT resources (in the Mist), and their combined interactions with Cloudlets and the Cloud. Cloudlets and Edge nodes are merged under one category, as they are both handled similarly in Cloud literature [1, 2].

control and data management to the edge of the network, far beyond the naive view of resource limitation. In the remainder of this section, we overview the major challenges in *IoT migration toward the edge*, especially in contrast to the dominant approach of centralized and proprietary IoT proliferation that is governing most solutions [10]. The major challenges witnessed in cloud IoT architectures are depicted in Fig. 2, wherein we annotate the span of fog networks, as well as the challenges under each tier in the hierarchical view of cloud components.

Spatial Correlation

Determining the location of collected data is becoming an increasing challenge in IoT systems. While advances in GPS as well as indoor localization systems have enabled sub-meter localization, many IoT nodes do not encompass the resources to self-localize. More importantly, many IoT systems are building on archaic localization schemes from wireless sensor networks (WSNs), which were largely static in deployment, or had specific mobility patterns that may not fit most IoT scenarios. As IoT applications are mandating better coupling of data generation and coverage accu-

While advances in GPS as well as indoor localization systems have enabled sub-meter localization, many IoT nodes do not encompass the resources to self-localize. More importantly, many IoT systems are building on archaic localization schemes from Wireless Sensor Networks, which were largely static in deployment, or had specific mobility patterns that may not fit most IoT scenarios.

As we witness most of today's things turning into micro-computers with communication and identification capabilities, the emphasis on uniform expression of functional capacities of IoT resources is growing in importance. That is, as we attempt to interoperate between IoT architectures, we need to have yardstick methods to identify and evaluate the functional capabilities of IoT resources across heterogeneous deployments.

racy, many schemes are challenged by improving the latter. This becomes more of a problem when data pruning and averaging techniques attempt to align and fuse sensed reports from IoT systems, which is further exacerbated by the heterogeneity of IoT devices.

There is promise in establishing localization in cloudlet zones, especially as they are intrinsically confined to pre-determined regions, and a "local-global" view of available localization schemes could be fused to improve the spatial correlation of data. Moreover, as singular IoT systems may fail to individually localize their data, or establish coverage in a given region, leveraging cloudlet knowledge of overlapping IoT deployments may increase the spatial knowledge of data from a given region.

Temporal Limitations and Service Latency

Real-time access to data sources is pivotal to sense-making systems in the IoT. Many of the proposed solutions for smart cities require high levels of coordination between IoT systems, and high responsiveness is a core mandate. The challenge of leveraging cloud resources is the inevitable queuing delay aggregated over multiple hops toward a cloud service, in addition to service time. Many experiments [1] have been carried out to demonstrate the challenge with latency in soliciting cloud resources.

On the other hand, bringing most IoT management closer to the edge yields significant interoperability challenges across IoT systems, in addition to lacking the infrastructure to mediate heterogeneous nodal operations. This is further worsened by the mis-coordination of communication between IoT nodes that not only differ in their duty cycling schemes, but also exhibit varying operation levels as per their power mandates and accessibility to their tethered devices.

Energy Footprint of IoT Operation

Most IoT systems are designed to conserve power in light of their individual operational mandates, so any attempts to interoperate IoT systems yields significant discrepancies in duty cycling schemes and multi-tiered operational levels. More importantly, to conserve power, most IoT nodes are designed to switch to low-profile sleep states to conserve power and are rarely open to IP-based probing from other Internet devices. While this is necessary for operational longevity, it affects interoperability across IoT systems. More critically, most of these sleep schedules are mandated by governing base stations and/or remote controllers, thereby limiting "visibility" of resources to neighboring IoT systems. As we attempt to merge traffic and data closer to the edge to conserve networking resources, a pressing challenge in IoT longevity will prove to be a hindrance.

On the other hand, recent research on the energy footprint of cloud architectures, especially as data centers are ever growing in their power demands [11], is offering new insights into the potential gain as we migrate IoT operation to the network edge under a broad view of fog networks. There is evident power gain in reducing overall network traffic, especially as we prune superfluous data before burdening cloud services up the hierarchy. Moreover, more contextualization of data, due to fog processing, may aid pruning and decision making that does not need to burden cloud systems.

The power footprint will likely dominate the *offloading granularity* problem. That is, deciding what should be processed at the user tier, what could be distributed on neighboring resources in the mist, what can be offloaded to context-aware cloudlets, and what demands high-power processing at the cloud is a major research challenge. These questions build on our collective expertise in elastic processing, network traffic engineering, big data management, and hierarchical fusion techniques.

Functional Mismatch under Cloud-Centralized Operation

IoT systems have long been developed as hard-coded architectures with pre-determined operational mandates. As we witness most of today's *things* turn into micro-computers with communication and identification capabilities, the emphasis on uniform expression of functional capacities of IoT resources is growing in importance. That is, as we attempt to interoperate between IoT architectures, we need to have yardstick methods to identify and evaluate the functional capabilities of IoT resources across heterogeneous deployments.

This is a precursor to enabling IoT cooperation at the network edge, as we attempt to leverage centralized service matching carried out by cloud services that do not have accurate or real-time feeds of which IoT nodes are currently duty cycling, offering their services, accessible in a given region, or reachable via a reliable networking medium [3]. The general assumption of cloud-IoT systems that simplify a global view of what is accessible, and carries out offline matching between service requests and actual IoT resources, is no longer feasible as our IoT systems grow ever more mobile and independent in operation [5].

Privacy and Security

Recent advancements in developing IoT-specific privacy and security mechanisms, such as the O Auth – 2.0 protocol, are tackling one of the most hindering factors in IoT traction. However, many of the challenges with IoT privacy and security result from the remote management of these important operations. For example, in data denaturing (e.g., blurring out the faces of pedestrians in a cloud-camera architecture) is often carried out at remote cloud services, presenting multiple opportunities for data breaches along the propagation links. Much of the context of the data is also lost in this cloud offloading approach, whereby important correlations between data from a single zone might be lost in the mass-scale processing of collected data. In addition, the challenge of data anonymization is magnified as more central modules have larger visibility to all data collected from multiple edge zones. For example, think of a FitBit server that observes all the movement and health patterns of all of its users, and then anonymizes the results prior to company-wide studies.

In terms of security, there are significant challenges in centralized "blanket" methods that are applied to securing data at cloud servers, or attempting to burden low-end IoT devices with encryption and authentication. While advances in the OAuth – 2.0 protocol are yielding promising solutions, much has to be done to enable end users in securing their own data, and deciding on the frequency and quality of data that is reported from their end devices to upper-tier cloud components.

Memoryless Operation

There is a rising challenge in maintaining user profiles in each zone, to establish factors such as trust in data reporting, weeding out false/malicious reports, and promoting "trusted" users in a given IoT application, especially in crowd-based scenarios. However, as users migrate from one zone to another, the exchange of this information between cloudlets opens up many security and privacy challenges, in addition to challenges in cooperation between heterogeneous cloudlet architectures. The scope of a designated zone, and what information about it and its ensuing users could be collected, exchanged, and analyzed, remains a significant challenge as we bring more processing and decision making power to the edge of the network.

Recent Advancements: Enabling Fog IoT

The abundance of smart devices in our everyday interactions is mandating novel approaches to viewing what the IoT encompasses, and the aggregated power of these resources. Recent developments in cloud computing are already promising many advancements in reliable service delivery via cloudlets [1], as well as resource provisioning on both the cloud and cloudlet levels.

Furthermore, most of today's smart devices are able to multi-home, whereby a typical smartphone can communicate over LTE, Bluetooth, WiFi, and ANT+ all in one device. This development is enabling many devices to act as mediators between multiple IoT systems, and further potentiates IoT interoperation in the mist, whereby neighboring nodes that have similar multi-homing capabilities can establish multiple overlay networks for different applications and/or services.

The development of nano data centers, building on the growing potential of smart devices and vehicles, will enable more data storage at the fog level. This will enable both rapid access to generated data and near-real-time probing of resource profiles in a given fog region. That is, a service can actively explore nearby fog resources (e.g., sensors) and query the data it collects; for example, to verify if there is a hit to the query, or if other resources should be solicited. More importantly, we can build on existing data aggregation techniques [12] that require localized data storage to enable better pruning and data management at the edge. These locations could coincide with cloudlets to bring more processing and intelligence to the network edge and reduce IoT traffic burdening the larger network backbone.

Recent advancements in short-range communication protocols, summarized in the table depicted in Fig. 2, are promising low-power and long-range communication between user devices, neighboring devices in the mist, as well as long-range communication with cloudlets. This is highly utilizable in scenarios where mobility-driven communication protocols, such as Dedicated Short-Range Communications (DSRC), are enabling cooperative operation in vehicular IoT systems and vehicular clouds. The scale is indeed ever expanding, from nano-communication with brain-machine interfaces via neuro-dust sensors [13], to long-range/high-bandwidth communication witnessed in the IEEE 802.11 family. In addition, recent developments in narrowband IoT (NB-IoT) and cloud radio access network (C-RAN)-based IoT developments [14] are expanding the scope of which IoT systems we can communicate with on the cellular backbone, with the added benefits of reliable and high-bandwidth channels.

Furthermore, we are witnessing the development of many resource discovery protocols that are enabling real-time probing and utilization of IoT resources. Whether this is carried out on the cloud, edge, or mist level, there is great potential in the mechanisms being developed to interrogate and register resources in real time, and scale their inclusion in fog-level resource pools, for service matching [7, 12].

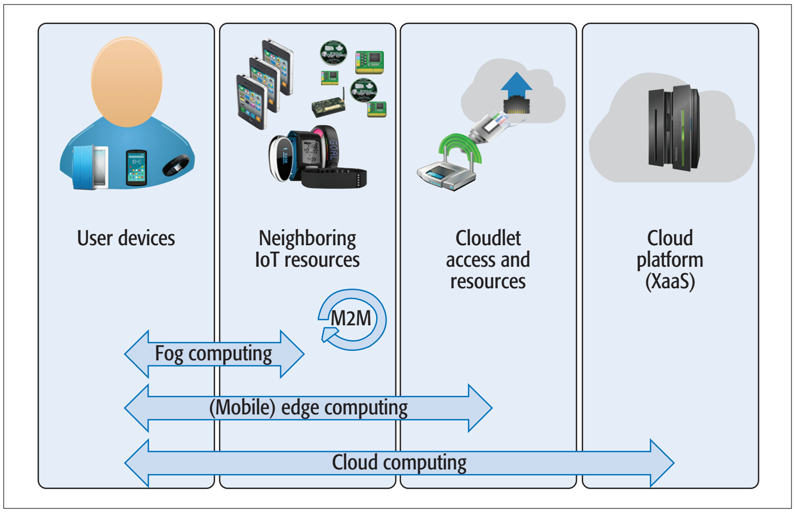

Figure 3. The interactions between a user-centric fog IoT architecture and cloud variants. The scope of fog IoT lies between user devices and neighboring IoT resources (i.e., the first two tiers). In a simple use case, user-centric health applications can probe local resources around the user (e.g., smart wristband, chest strap, blood pressure monitor) and correlate with nearby resources (e.g., nearby temperature sensors and weather stations) to establish whether certain readings (e.g., higher heart rates) could be influenced by ambient weather conditions or potentially a user-specific condition. More thorough analysis of history logs and user data could be accessed via cloudlets and edge computing. Finally, epidemic monitoring and crowd-level analytics could be offloaded to remote (and more powerful) cloud services over their multiplicity of XaaS services.

Toward a Fog-IoT Architecture

As we advocate for moving from a service-centric (cloud/edge) to a user-centric (fog) approach to IoT systems for smart cities, we focus on the architectural components that will enable such a progressive framework. At its core, a user-centric architecture must utilize the context as well as resources of a local fog, and establish real-time management modules that will tap into the potential of neighboring resources in the mist, as well as cloudlet/edge-level resources when needed. Thus, service matching, mobility monitoring, and overall offloading granularity are largely served within the bounds of the fog network rather than the coud.

On an architectural level, we advocate for establishing an IoT-in-the-fog controller that is able to probe local resources and communicate directly with a local fog mediator, which could be the cloudlet/edge access point. The controller operation could be deployed on a dedicated device placed for that purpose (e.g., in a roadside unit) or delegated to high-end resources (e.g., a smartphone or IoT hub).

The core operational mandate of this controller would be to respond to policies mandated by the fog mediator, as passed down from respective cloud services, but matching the current resources in the fog zone. This includes catering to mobility and resource volatility, especially in utilizing mobile/vehicular resources in urban environments. Figure 3 overviews the interactions between cloud variants and what they are dubbed in current literature, highlighting the reach/scale of each cloud variant. A simple scenario for e-health applications is presented in Fig. 3, whereby classes of e-health applications running on each tier of the fog-IoT architecture are overlaid and explained.

However, what makes this architecture unique is that cloud-based IoT architectures are almost always service-centric (over the cloud/edge). However, the fog IoT architecture is envisioned to be user-centric, whereby interactions between devices, exchange of control messages, and data flow are governed by user-centric policies. For example, a user decides on the granularity of services and data they wish to access, and the associated monetary and energy cost of probing the provisioned resources.

While this entails more processing and power at the edge, it builds on many advantages in privacy preserving mechanisms, mobility control, and elastic offloading when the need arises. The granularity of data handled by users could further be controlled from both the user (to reduce access latency) and from the cloud (to enforce access rights).

High-Yield Research Directions

Memory-Preserving Operation in the Fog

As we attempt to enable smart city applications, there is a major opportunity loss in our *memoryless* view of contributing IoT systems. Within a given region, on a fog network scale, there is much that can be inferred and stored about region and user profiles per zone. That is, we can establish history-based logging of data and contributing users, adding to the development of trust-ranking schemes for each user known to commonly access IoT systems in a given zone.

Furthermore, there is great promise in establishing time-series-based inference of potential resource needs in a given region, based on maintaining *memory* of what is being produced (in terms of data) and accessed (in terms of IoT resources) in a given region.

In Sensing Architectures

We advocate for a migration from sensing as an intrinsically event/sampling-based paradigm to a service paradigm. That is, data is only collected when there is a demand for a given service, and the decision of which nodes in an IoT ecosystem are to take part in sensing (i.e., load balancing) should be made on a fog level rather than an individual IoT system level. While many applications will mandate that their own sensors report data (e.g., for reliability and calibration constraints), there is significant data redundancy across IoT systems, which is causing significant big sensed data challenges in IoT scalability and management.

Building on ICNs

There is great promise in the recent development of information-centric networks (ICNs) that handle data at an intrinsic network primitive. In ICNs, data is automatically encoded and distributed over the network architecture, masking many of the challenges of data naming and cloud-based querying over dedicated IoT systems. However, much investigation is needed in terms of enabling remote "subscriptions" to data from given IoT systems, to enable IoT nodes to act as data/content providers in an interactive environment that responds in real time to demand and popularity metrics, rather than collect data in the hope of future interest.

Incentive Schemes and Interplay between Cloud Variants

A major challenge in crowd-based IoT systems is soliciting data and resources from users. Many recent research endeavors have investigated incentive schemes that address this challenge and the promise of these systems [15] in yielding higher user contributions. However, much is to be discovered in incentivization across IoT platforms in smart city environments. This includes how to establish incentives across IoT systems to solicit the best-fit resources for a given task in a market-driven architecture that reacts to resource abundance, and responds to urgency in service matching and timely delivery of IoT services.

Much is to be discovered in incentivization across IoT platforms in smart city environments. This includes how to establish incentives across IoT systems to solicit the best-fit resources for a given task in a market-driven architecture that reacts to resource abundance, and responds to urgency in service matching and timely delivery of IoT services.

Interplay of IoT Services and Service Orchestration Platforms

One of the great promises of IoT is the potential to build larger services (e.g., weather prediction and route planning) based on atomic/simpler services (e.g., temperature sensors and road monitoring cameras). The premise of service orchestration hinges on the accessibility of reliable services that are closer to the edge, with capped access latencies, and contextually enforced data collection and pruning mechanisms. That is, we need to develop more robust and reliable atomic services to feed larger service orchestration platforms. The fog IoT architecture can synergize heterogeneous services and architectures at the edge level, and developments in policy management, data pruning, and information dissemination at fog services will potentiate service orchestration.

Concluding Remarks

The potential of IoT proliferation in the fog is evident in the development of many technologies that bring more resources to the network edge. For over 20 years, sensing systems and IoT were envisioned as technologies that require *light* operation at the edge, with emphasis on cyber foraging and cloud offloading to enable reliable services. The rise of mobile edge computing, cloudlet access, and M2M communication modes are all providing ample resources for migrating more processing and resource management at the network edge. In this article we survey many of the challenges in attempting to remotely operate sensing systems, and the ensuing big sensed data challenges that warn us of data generation beyond what we can communicate and process. As more researchers are advocating for migrating IoT data management to the network edge, utilizing variants of cloud computing paradigms, this article surveys the core challenges in this migration and proposes a roadmap for IoT interactions on the fog/edge/cloud tiers, based on the aforementioned developments in edge technologies. Finally, we present a number of high-yield directions that will further propagate IoT development in the fog. It is important to note that many developments are taking place in parallel research domains, and it is at the heart of this article to highlight the potential benefits in synergizing some of these mainstream efforts. This has been surveyed in Fig. 2 as a building block to instigate further discussions on cross-domain synergy toward more potent fog IoT architectures.

For over 20 years, sensing systems and IoT were envisioned as technologies that require light operation at the edge, with emphasis on cyber foraging and Cloud offloading to enable reliable services. The rise of Mobile Edge Computing, Cloudlet access, and M2M communication modes are all providing ample resources for migrating more processing and resource management at the network edge.

References

[1] M. Satyanarayanan, "The Emergence of Edge Computing," *IEEE Computer*, vol. 50, no. 1, Jan. 2017, pp. 30–39.

[2] S. Sarkar, S. Chatterjee, and S. Misra, "Assessment of the Suitability of Fog Computing in the Context of Internet of Things," *IEEE Trans. Cloud Computing*, Oct. 2015, pp. 1–14.

[3] A. Al-Fuqaha *et al.*, "Toward Better Horizontal Integration among IoT Services," *IEEE Commun. Mag.*, vol. 53, no. 9, Sept. 2015, pp. 72–79.

[4] F. Bonomi *et al.*, "Fog Computing and Its Role in the Internet of Things," *ACM Wksp. Mobile Cloud Computing*, Aug. 2012, pp. 13–16.

[5] S. Oteafy and H. Hassanein, *Dynamic Wireless Sensor Networks*, Wiley-ISTE, June 2014. ISBN: 978-1-84821-531-3.

[6] J. Preden *et al.*, "The Benefits of Self-Awareness and Attention in Fog and Mist Computing," *Computer*, vol. 48, no. 7, July 2015, pp. 37–45.

[7] S. Oteafy, "A Framework for Heterogeneous Sensing in Big Sensed Data," *IEEE GLOBECOM*, Dec. 2016, pp. 1–6.

[8] M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *IEEE Personal Commun.*, vol. 8, no. 4, Aug. 2001, pp. 10–17.

[9] S. Oteafy and H. Hassanein, "Resilient IoT Architectures over Dynamic Sensor Networks with Adaptive Components," *IEEE Internet of Things J.*, vol. 4, no. 2, Apr. 2017, pp. 474–83.

[10] A. Al-Fuqaha *et al.*, "Internet of Things: A Survey on Enabling Technologies, Protocols, and Applications," *IEEE Commun. Surveys & Tutorials*, vol. 17, no. 4, 2015, pp. 2347–76.

[11] F. Jalali *et al.*, "Fog Computing May Help to Save Energy in Cloud Computing," *IEEE JSAC*, vol. 34, no. 5, May 2016, pp. 1728–39.

[12] J. Sahoo, S. Cherkaoui, and A. Hafid, "Optimal Selection of Aggregation Locations for Urban Sensing," *Proc. 2014 IEEE IC)*, Sydney, Australia, Aug. 2014, pp. 239–44.

[13] M. Maharbiz *et al.*, "Reliable Next-Generation Cortical Interfaces for Chronic Brain-Machine Interfaces and Neuroscience," *Proc. IEEE*, vol. 105, no. 1, Jan. 2017, pp. 73–82.

[14] A. Radwan et al., "Low-Cost On-Demand C-RAN Based Mobile Small-Cells," *IEEE Access*, vol. 4, May 2016, pp. 2331–39.

[15] D. He, S. Chan, and M. Guizani, "Privacy and Incentive Mechanisms in People-Centric Sensing Networks," *IEEE Commun. Mag.*, vol. 53, no. 10, Oct. 2015, pp. 200–06.

Biographies

Sharief M. A. Oteafy [M] is an assistant professor at the School of Computing, DePaul University. He is actively engaged in the IEEE Communications Society, serving as the ComSoc AHSN Standards Liaison and on the ComSoc Tactile Internet Standards WG. He co-authored a book, *Dynamic Wireless Sensor Networks* (Wiley), and has presented over 50 peer-refereed publications and delivered multiple ComSoc tutorials on sensing systems and IoT. He has co-chaired a number of IEEE workshops, in conjunction with IEEE ICC and IEEE LCN. He is an Associate Editor of *IEEE Access* and on the Editorial Board of Wiley's *Internet Technology Letters*.

Hossam S. Hassanein [F] is a professor and director of the School of Computing at Queen's University, Kingston, Ontario, Canada. His record spans more than 500 publications, in addition to numerous keynotes and plenary talks at flagship venues. He is also the founder and director of the Telecommunications Research Lab at Queen's University. He is an IEEE Communications Society Distinguished Speaker (Distinguished Lecturer 2008–2010). He is the past Chair of the IEEE Communication Society Technical Committee on Ad Hoc and Sensor Networks. He has received several recognitions and best paper awards at top international conferences, and led a number of symposia at flagship ComSoc conferences.