

Demystifying Fog Computing: Characterizing Architectures, Applications and Abstractions

Prateeksha Varshney and Yogesh Simmhan

Department of Computational and Data Sciences, Indian Institute of Science (IISc), Bangalore 560012, India

Email: prateeksha@grads.cds.iisc.ac.in, simmhan@cds.iisc.ac.in

Abstract—Internet of Things (IoT) has accelerated the deployment of millions of sensors at the edge of the network, through Smart City infrastructure and lifestyle devices. Cloud computing platforms are often tasked with handling these large volumes and fast streams of data from the edge. Recently, Fog computing has emerged as a concept for low-latency and resource-rich processing of these observation streams, to complement Edge and Cloud computing. In this paper, we review various dimensions of system architecture, application characteristics and platform abstractions that are manifest in this Edge, Fog and Cloud ecosystem. We highlight novel capabilities of the Edge and Fog layers, such as physical and application mobility, privacy sensitivity, and a nascent runtime environment. IoT application case studies based on first-hand experiences across diverse domains drive this categorization. We also highlight the gap between the potential and the reality of Fog computing, and identify challenges that need to be overcome for the solution to be sustainable. Taken together, our article can help platform and application developers bridge the gap that remains in making Fog computing viable.

I. INTRODUCTION

The prime drivers of Big Data over the past decade have been the WWW, social media, eCommerce, and enterprise information systems, with data and services being consolidated in public and private data centers. However, *sensing at the edge* of the network is accelerating many-fold. This is enabled through infrastructure like Smart Cities that are being deployed globally, lifestyle devices like wearables and smart appliances, and observations collected *ad hoc* through crowd-sourcing over smart phones [1], [2]. A large class of this data streaming from the edge can be attributed to the *Internet of Things (IoT)*, where the convergence of commodity sensors, accessible communications, and the need for intelligent infrastructure management is driving this exponential growth.

The traditional model of acquiring and processing this data from the edge to offer useful services has been through *Cloud computing*. These large Cloud data centers host computing and storage infrastructure with high bandwidth connectivity, and allow IoT applications and services to be hosted remotely [3]. Their on-demand access to seemingly infinite resources accessed through simple web service APIs, combined with the pay-as-you-go billing model that capitalizes on economies of scale have made Cloud computing popular.

At the same time, two key challenges emerge with this hub-and-spoke model. Public Cloud data centers are globally distributed, but typically limited to one or a few per country (depending on the size). This means that the round trip latency (depending on the size). This means that the round trip latency (RTT) from the edge to the Cloud is $\mathcal{O}(100)$ ms [4]. Many IoT applications are *latency sensitive*, like demand

prediction for Smart Power Grids [5], [6] and voice response in Apple Siri. Hence, hosting the analytics and decision making in the Cloud can compromise their performance [7], [8].

Similarly, video is emerging as a major source of data from the edge, from urban surveillance and ATM cameras that are part of the infrastructure, to personal body cameras and crowd-sourced recordings from phones [1], [9]. Here, the *bandwidth* required to push high fidelity video streams to the Cloud is prohibitive, and this further increases the RTT if the video analytics on the Cloud has to control edge devices, like zooming a PTZ Camera or detecting faces in real-time.

This has led to a design paradigm of processing at the edge to supplement, or replace, Cloud computing [10], [11]. Also called *Edge Computing* or *Mobile Cloud* [3], [9], these are typically used in vertically integrated applications where a part of the processing and analytics happens on the edge device while the Cloud is used for coordination and data archival. E.g., a FitBit fitness watch may pre-process and visualize data on the smart phone before archiving to the Cloud.

However, Edge computing suffers from several deficiencies. The edge platforms tend to be constrained devices, with battery capacity or memory often being the limiting factor rather than even compute capability. This can cause resource contention if multiple IoT applications need to be deployed concurrently [7]. Also, there are no well-defined or robust runtime or management platforms for composing generic Edge computing applications, comparable to the virtualization or service-based architectures exposed in Cloud computing. As a result, edge computing is limited to bespoke solutions.

Fog computing [8], also known as *Cloudlets* [4], was introduced by Satyanarayanan, et al., and popularized by Cisco as a complementary resource-rich layer that sits between the edge device and the Cloud. [4] suggest that this layer be one network hop away from the edge to offer both low and predictable latencies to support gaming and video conference services. The Cloudlets are meant to serve as small-scale data centers that are placed closer to the edge [10], where often both the generators of the observations and the consumers of the analytics reside. At the same time, the Fog layer can offer easier manageability of resources as services, and possibly a feasible business model similar to the popular Clouds.

In this regard, Fog computing has similarities with *Content Distribution Network (CDN)*, that sits close to the edge for low latency delivery of content, typically static or slow changing and populated from the Cloud, but are otherwise passive hosts. [1] considers Cloudlets as a “CDN in reverse”, with the edge populating the Cloudlet, and the Cloudlet serving data to the

Cloud. At the same time, the Fog can also actively host services and analytics – closer to the edge than, even, CDN.

From a societal perspective, Fog computing also becomes relevant when there are network outages that restrict or remove the ability to reach the Cloud data center. E.g., the *Cyclone Vardah* that hit Chennai, India in Dec. 2016 damaged a trans-oceanic optical cable connecting India with Singapore and Europe, where many commercial data centers are located¹. In future, as critical services such as water, power and transport are controlled by smart algorithms, hosting them in a Fog layer with peering mechanism allows a graceful degradation of functionality even in the absence of access to the Cloud, rather than cause a debilitating loss of city services.

However, it is important neither to over-state the potential of Fog computing nor to dismiss it as hype. Despite the existence of the concept for several years now, commercial deployments of Fog Computing are yet to take off. Part of the reason is that there are still an inadequate number of widely deployed or critical applications that find the Fog to be essential. There is also a lack of clarity on the application model, runtime and management environments for a Fog platform. And lastly, a sustainable business model and service providers are still evolving. So it is useful to take a *critical look* at Fog computing as a *prospective* pervasive platform.

In this paper, we attempt to characterize some of the key features of the system design and computing architecture of Fog computing; the application features, particularly from IoT use-cases, that motivate the need for Fog computing; and lastly the programming models and abstractions that can be leveraged to bridge the gap between the applications and the systems for designing intuitive runtime platforms.

Several papers offer desiderata for Fog computing and its concepts. Many tend to be superficial or narrowly defined. Cisco subsumes the Fog into the Edge (or vice versa), there by discounting widely-available devices like smart phones and Raspberry Pis deployed as part of Smart City infrastructure [8]. [12] attempts to define Fog computing, but views it from the limited prism of network management and connectivity. Several papers on Cloudlets offer exemplar applications that are actually deployed, but fail to offer an overarching technical view of this space [1], [4], [13]. Others propose a programming model for composing applications that run across mobile, Fog and Cloud layers as a Platform as a Service (PaaS), but limit it to a rigid hierarchy [14]. [15] attempts a similar effort as our paper, but favor mobile devices rather than edge devices at large. The latter, many of which can be part of an IoT infrastructure, are key. [7] discuss the role of Fog computing on IoT applications, and highlight open issues. These are further discussed in the related work section (§ V).

Here, however, we go a step further and provide intrinsic dimensions and capabilities of Fog computing that should be of interest to platform developers and application designers, while recognizing that this is still an emerging space. These features distinguish the Fog from traditional mobile or Cloud computing but *form a continuum*. These characteristics are as yet inadequately studied in the context of a holistic Edge, Fog

¹A huge storm has messed up India's Internet by Rishi Iyengar, CNN, December 14, 2016 <http://money.cnn.com/2016/12/14/technology/india-cyclone-varDAH-chennai-internet/>

and Cloud computing eco-system. We also motivate the problem space with emerging, rather than futuristic, applications.

We make the following specific contributions in this paper:

- 1) We characterize the *expected capabilities of a Fog computing system*, and lay them in the context of existing Edge/Mobile and Cloud computing architectures. We highlight dimensions that distinguish these alternative, but inter-linked, resource layers (§ II).
- 2) We review the *application requirements* that motivate the need for a Fog layer, and identify the gaps posed by a Cloud-only or Edge+Cloud model. We also offer case studies of applications and their features (§ III).
- 3) We highlight the possible *runtime and middleware capabilities, and application models* required in a Fog layer, while recognizing that the Fog platform will need to adapt to the unique needs of applications and the actual manifestation of the Fog system (§ IV).
- 4) Lastly, we *discuss* the gap between potential and reality of Fog computing, and identify challenges that need to be overcome and solutions that are possible in this evolutionary phase (§ VI).

II. FOG COMPUTING SYSTEM ARCHITECTURE

A. Definitions

There are various overlapping definitions of Fog computing. We summarize some of these views next, but then offer our detailed characterization rather than just a simplistic definition.

- Satyanarayanan, et al, [16] state that Cloudlets, synonymous with Fog computing [17], are a resource-rich (cluster of) computers that are located one hop away from the mobile devices. They have Gigabit interconnect, and high bandwidth, through Wireless to the mobile device and through WAN to the Internet. The Cloudlets offer a “data center in a box” close to the mobile device in order to reduce the network latency and bandwidth-induced latency to support interactive applications.
- Cisco [8] views Fog Computing as offering resources like compute, storage and networking similar to Clouds, with support for virtualization and multi-tenancy. These are geo-distributed to offer low and predictable latencies to client applications that are mobile or part of the infrastructure. They do conceive of the Fog as having modest resource capabilities, ranging from an edge network router to a high-end server.
- [12] define Fog computing as a large collection of heterogeneous and decentralized devices, communicating among themselves to store data and process tasks, that can be leased to users to support basic network functions or sandboxed applications. They take a communication-centric view of Fog computing, with 4G/5G connectivity enabling the Fog layer and Software Defined Network (SDN) management pushed to the Fog.

B. Comparing Edge, Fog and Cloud Computing

From the above definitions, among many others, Fog computing is seen as a resource layer that fits between the edge devices and the Cloud data centers, with features that may resemble either. Hence, it is worth comparing and contrasting




Figure 1. Resource characteristics of Cloud, Fog and Edge computing systems

the characteristics of the Fog computing system architecture relative to the capabilities of Edge and Cloud computing.

1) *Resource Characteristics*: There are several resource and performance characteristics that distinguish these three layers, as captured in a pyramid structure in Fig. 1. Chief among them are the *network characteristics*. The Fog is close to the edge in the network topology. Hence, it has a lower latency for access in *both* directions, i.e., serving content from Fog to Edge, and pushing data from Edge to Fog [16]. Whether it is at a 1-hop or multi-hop depends on the deployment. Further, the bandwidth between the Edge and Fog is also higher (e.g., using WiFi) than from the Edge to the Cloud (e.g., using cell networks). The Fog is also expected to have a reliable and high bandwidth Internet link [18], but its connectivity to the Edge may be less robust due to the use of wireless links for the last mile [15].

In these regards, the Fog layer is analogous to a CDN but typically closer to the Edge [12]. The *storage capacity and data longevity* of the Fog layer is much higher than the Edge devices, though more limited than Clouds. This storage can be used to cache large datasets that are useful to edge devices, such as apps or Virtual Machine (VM) images that have to be installed on devices for Smart City infrastructure, or snapshots of historic data [10]. But more interestingly, the Fog also serves as a “reverse CDN” to allow edge devices to push data to the Cloud [19]. This allows scenarios where data is staged in the Fog and periodically pushed to the Cloud for archival, potentially after some pre-processing such as deduplication [1], [20]. Further, edge devices can also access data pushed up to the Fog by other edge devices [15].

At the same time, the Fog layer also offers *compute resources* with a higher capacity and reliability than the Edge but to a smaller scale than Clouds [8], [21]. This capacity can be used to host applications and services that range from Video Analytics as a service for processing video frames close to the Edge or directory services where devices register and access location-sensitive configuration for mobile phones [22].

The sheer number of Fog appliances will dwarf the number of Cloud data centers (even if not their cumulative compute power), just as the number of edge devices number in the billions. This makes *manageability* of the Fog more challenging, but easier than the Edge. The economies of scales will also come into play if Fog hardware is standardized, similar to commodity smart phones or shipping containers with Cloud hardware. This can make Fog computing affordable and (if the edge device is not captive) cheaper than Edge computing, though Clouds will retain their *pricing* advantage.

Lastly, the *energy profile* can influence the capability and availability of some resources. Edge devices are often concerned with battery life, and the choice of using specific Edge features can depend on current battery level [23]. Cloud data




Figure 2. Physical presence vs access of Cloud, Fog and Edge resources

centers reduce their energy footprint, but to limit operational costs [10]. The Fog is expected to run off grid power and, like the Cloud, be energy conscious to lower pricing [12], [20], [22]. But there may be remote places where the Fog runs on renewables like solar, when energy-aware usage is a key goal.

2) *Physical Presence and Access*: Besides network distance, it is also worth considering the physical distance between the three computing paradigms, and their accessibility by clients. In Fig. 2, four quadrants are formed from considering whether the resources within a layer are *physically centralized or distributed* (Y Axis), and whether their *access is global or restricted* (X Axis).

Resources in a Cloud data center are centrally located, but depending on whether the Cloud is public or private, are available to anyone in a pay-as-you-go model, or only to users of the private corporation [15]. That said, public Cloud providers host geographically distributed data centers, sometimes several in a country or continent, while the number of large private data centers for an enterprise is more limited.

We can similarly compare Edge and Fog resource distribution and access. Edge resources such as smart phones and set top boxes are distributed far and wide, but their access is restricted to individual users or managed applications [8], [9]. Fog resources are also physically distributed to be close to the edge, but not as dispersed. Additional specializations, on whether there is a Fog for each city block, one for the whole city or other variants, depend on the business models and applications that will evolve. One also expects the Fog to offer as a shared, pay-as-you-go IaaS or PaaS model [9], [10].

The distributed nature of Edge and Fog resources increases their *attack and failure surface*. There is a high chance of some device or Fog server failing or a network link dropping, and resiliency has to be built into the platform and application if necessary [24]. Clouds being more centralized are single points of failures, but Cloud fabrics and Big Data platforms internalize faults within the data center. But, as mentioned, natural disasters can cause these central hubs to get islanded, and bugs and security breaches can cause massive data loss.

The access restrictions on private Clouds and Edge devices translates to a *zone of trust* for applications and services hosted on them, which enables sensitive data and services to be hosted on them. Fog and public Clouds, however, are designed as shared resources with *multi-tenancy*, which require higher measures of security and sand-boxing between different applications or users, using containers or hypervisors [7], [12].

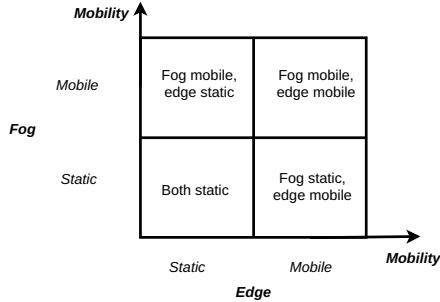


Figure 3. Physical mobility of Fog and Edge resources

That said, there may be Fog architectures where the resources are deployed for specific applications or organization (e.g., a Smart City municipality) and not available for rent, similar to a private Cloud [15]. Further, the Fog may sit at the boundary between public and private networks. E.g., in our IISc Smart Campus project ², hundreds of edge devices and sensors connected to the campus LAN monitor the water and power usage, and Fog servers on the DMZ connect to both the private LAN and the public WAN. So the Fog is well-suited to run proxy services that translate from one zone of trust to another, one service layer to another (e.g., CoAP to HTTP), or one network protocol to another (e.g., IPv6 to IPv4).

3) *Mobility*: It helps to understand the impact of mobility on these three resource layers as this impacts the communications, applications and platform design. We distinguish between mobility of the *physical resource*, discussed here, and mobility of the logical applications, which we examine later.

Cloud data centers, obviously, are not mobile though their platforms can ease the mobility of data and applications among their data centers. However, spatial mobility at edge devices is frequent, though not universal, e.g., mobility is seen in ubiquitous smart phones and autonomous vehicles, while they remain static in traffic cameras and smart power meters [21], [22]. Likewise, the Fog layer can also be manifest as a static or mobile resource [15]. A Fog server can be installed at fixed sites such as a coffee shop or the airport, or on mobile vehicles such as taxi cabs or trains. Fig. 3 captures these scenarios.

This mobility can have consequences on the link between the Edge and the Fog, or from the Edge to the public or private network, and these depend on the communication protocols used. However, we need to recognize that mobility can cause the network connectivity to be intermittent [25], causing a transient loss of data or compute access between the Edge, the Fog and the Cloud. But, in technologies where the access is based on physical proximity, such as Bluetooth, Near Field Communication (NFC) or line-of-sight networks like Millimeter Wave, the disconnection can be permanent [12].

While we do not expect this ephemeral behavior for links between the Fog and the Cloud, such cases are possible in interaction among edge devices or with the Fog. E.g., we use the notion of a “Data Sherpa” [23] to transfer observed data from a static sensor (edge) to a mobile smart phone (edge) using Bluetooth when they are close by, and then use the smart phone to push the data to the Fog or Cloud. But the

interaction between the two edges is opportunistic, not planned or repeatable [25]. Similarly, users in a train can use their personal devices with the Fog servers in the coach to access media or make use of a high-speed Internet, but only while they are riding. Connected cars (Fog) can peer with nearby vehicles to share information on road condition and collaboratively navigate traffic [15]. As a result, depending on the mobility of these two layers, the application and platform will need to be designed based on *permanent, transient, periodic, or ephemeral connectivities* between the layers and within the layers which can determine the *reliability of access* to data, storage, network and computing resources.

Lastly, applications may also relate spatial proximity with *trust and context*. Fog layers physically close to the edge may be trusted by the edge applications to share sensitive data or for the Fog to offer services. This is already used to start cars using a “key fob” present in the car rather than a key in the ignition, or using NFC to share contacts or Bluetooth to pair devices. This can be leveraged more actively by applications for trusted interactions between the Edge and Fog layers. This closeness also offers context for the interaction and can be used to serve content relevant to the person from the Fog [9].

III. APPLICATIONS FOR FOG COMPUTING

In this section, we go in-depth into the characteristics and requirements of applications that are well-suited to different forms of Fog resources, in concert with the Edge and Cloud resources. We also offer detailed use cases on Fog applications.

A. Mobility of Data and Analytics

Similar to physical mobility of the resources, we can have logical movement of applications between Edge, Fog and Cloud layers. The entity that moves can be a *process* (e.g., analytics, services), *data* (e.g., video streams, device images) or *event* (e.g., control signals, emergency notifications). Each have their own requirements from the resources.

Processes require appropriate execution environments in the resource layer and adequate compute and memory capacity. While virtualized Clouds offer various VM images and sizes, Fog is expected to have a similar flexibility using containers or hypervisors [3], [20]. Edge may be the most constrained and have limited choices. Data movement requires bandwidth and storage to be available between the layers, and latency may be a concern too. These sizes range from Kilobytes for configuration files to Gigabytes for high-fidelity observations being archived [9]. Events streams are special as they are lightweight, and serve as control signals that prompt an action. They typically need low latency transfer from source to sink, and the ability for the receiving layer to react [21].

There are four major *logical mobility models* for the source, the sink, and their direction for this movement, as shown in Fig. 4. The first shows the logical entity move from Edge to Fog to Cloud, such as moving observation streams from Edge to Cloud for archival, with the Fog performing filtering or pre-processing [19]. The second shows the entity move from Cloud to Edge, e.g., flashing a new device image onto a Pi deployed in the field, with the image cached in the Fog. Unlike these two “open loop” models, IoT applications often have a “closed loop” that consists of Observe, Decide and Act, and

²IISc Smart Campus Project, <http://smartx.cds.iisc.ac.in>




Figure 4. Logical mobility of applications between Edge, Fog and Cloud layers. Arrows indicate direction, diamonds indicate optional processing or staging at that layer. Process, Data and Events can move in both directions, with any of the layers as source or sink.

the last two modes indicate these feedback cycles. E.g., the data source can be on the Edge, decisions are made at the Fog and/or Cloud, and the notification arrives back at the Edge for a response [3], [19]. Similarly, the Cloud could provide a data stream to the Fog or Edge that respond back with a control message to enact a service call in the Cloud. Of course, subsets of these paths may be relevant for an application as well, e.g., start at the Fog, move to Cloud and then to Edge.

B. Quality of Service (QoS) Characteristics

There are several application QoS requirements that can drive the mix of resource layers to be considered. *Latency and bandwidth* are two obvious factors for applications, and [26] has earlier grouped applications based on their ability to withstand such delay. *Inelastic applications* require real-time processing (i.e., hard latency constraints), typically related to safety of humans such as autonomous vehicles, or for applications consuming transient sensor streams [8]. The latency accumulates in each layer that the data passes through, and the Fog may offer better pair-wise cumulative latency with the Edge and the Cloud, than from Edge to Cloud. *Elastic applications* (different from Cloud elasticity) have flexibility in latencies and do not require immediate processing. These are delay-tolerant and designed for off-line or batch processing, such as surveys using drones for planning [22].

The *robustness* of the application becomes important as well. Mission-critical applications that operate City infrastructure such as power grids or traffic signaling may require deterministic network connectivity and computing capacity, even during disasters, that may be better provided close to the Edge or Fog [3], [21]. On the other hand, if the robustness relies on instantaneous availability of large computing capacity or scalable services, say for supporting periodic events at stadiums, Clouds offer this flexibility provided redundancy of communication channels and data centers are built in.

Lastly, *cost and user base* can play a major factor. For applications with a growing user community, the elastic scaling of the Cloud or the Fog allows them to pay proportional to their community size and revenue [10], [16]. Sometimes, applications coupled to consumer hardware, like autonomous cars and wearables, come with in-built (free) computing capability, on the device or a smart phone, that can (or should) be leveraged. Others like city utilities with a slow-changing population may wish to have captive resources, self-managed at the Edge and Fog, due to the scale that can be leveraged, and complemented by the Cloud [3]. Sometimes, market forces may increase the resource availability based on the demand

from these applications, such as Cloud data centers or cell phone towers coming up in regions of revenue growth.

C. Policies and Constraints

Privacy and security is a growing concern as the number of connected systems grow. There may be explicit corporate policies or national laws that do not allow data or services to be moved outside network or geographical boundaries [3]. Here, the Edge and the Fog may be critical to meet performance needs and legal compliance. But Cloud data centers in general have stronger physical and digital security systems that often surpass what can be provided at widely distributed Edge and Fog devices that may operate in hostile environments [13].

IoT applications are often designed as *Cyber Physical Social Systems (CPSS)*, that span data and software (cyber), infrastructure and devices (physical), and human interaction (social). In some cases, the application is intrinsically tied to specific devices (e.g., drones) [22]. In others, they may have explicit dependencies on computing infrastructure or resource capabilities. E.g., a video classification may require model-training to take place on GPUs or FPGAs, which may be available only on the Edge or Fog. The application may also be coupled to a specific mobile user, and it will need to use logical mobility to move state and services to the resource physically closest to the user [10]. So some of the applications may be constrained to those cyber and physical systems.

D. Application Use Cases

1) *Urban Surveillance*: There is a global push toward Smart Cities and India is poised to upgrade over 100 cities as part of a national initiative [23]. As part of a *Light-Pole Computing* pilot, several city blocks in Bangalore are installing video cameras, environmental sensors, Edge computing platforms (e.g., Raspberry Pi) and, at each city block, accelerated Fog computing platforms (e.g., NVIDIA Jetson TX1) [27].

Given the advances in deep learning, video surveillance in urban environments is essential not just for public safety but also as a proxy for ambient observations. E.g., these analytics can be used to automatically identify parking violations, classify vehicles and people, generate visual summaries for safety agencies, and even blur regions for privacy [9], [26]. Training the neural network models is computationally costly, and the source video streams at the edge are also large in size. Training may even be required every few hours to adapt to local conditions. Classification using trained models is faster, but we have observed even GPU-accelerated being able to process only hundreds of frames per second.

Here, the training can be done in batch but moving the compute and data to the Fog saves bandwidth. The classification can move to the Edge or Fog, based on latency needs. At night, when human safety is a concern, real-time needs trump a daytime use case of detecting parking violations. As an example of a feedback control loop, the PTZ camera can be controlled by the analytics to automatically zoom into features of vehicles. Safety situations detected at the Edge can trigger an alert to spatially proximate smart phones, or notify officers through the Cloud [14]. The Edge and Fog platforms can also host other IoT applications from public utilities that tap into the video or sensor streams for event-analytics [11].

2) *Smart Power Grid*: The *Los Angeles Smart Grid* will serve over 4 *Million* customers in the largest public utility in the US [28]. Net-connected smart meters observe power demand at households and industries and report them periodically to the utility every few minutes. These run off 2G or P2P communication models. Further, Supervisory Control And Data Acquisition (SCADA) systems at hundreds of city feeders collect high-frequency data on voltage and current quality across multiple phases, often at KHz rates [29].

Demand-response (DR) optimization helps shape or shift power demand to match supply capacity [5]. Here, liability and privacy concerns mean that demand-response and load control decisions happen at multiple-levels. The utility uses global but coarse-grained data to run demand forecasts and determine the curtailment level required [30]. It then notifies the customer gateway, say an Edge device for a consumer or a Fog device for a campus like USC, of the curtailment required and the discount (or penalty) pricing. These gateways then take local decisions to determine curtailment strategies and control, say a smart appliance or an electric car, or centrally change set-points of HVAC systems across campus buildings.

While the DR feedback loop can tolerate delays of several seconds or minutes, *state estimation* to determine the health of the distribution network is time sensitive. Here, computational models run over fine-grained power quality measurements from feeders to detect instability, and take rapid response on the Edge to avoid the whole network from being affected. There are distributed state estimation models that are well suited for Fog devices, with periodic updates pushed to the Cloud and global aggregates pushed back. As one can imagine, cyber-security concerns of such critical infrastructure eliminate the possibility of having multi-tenancy on these Fog resources.

3) *Drones for Asset Monitoring*: Lastly, drones offer a novel and emerging platform that highlight the role of mobility in Edge and Fog computing. Swarms of these autonomous platforms are being put to use for asset monitoring of gas pipelines or city infrastructure. Drones have on-board navigation and environmental sensors, visual/IR cameras, and sometimes even actuators (e.g., probes, grapples) [22]. They have wireless tethers, often on a line-of-sight, with mobile or static base-stations that act as the Fog layer. Upon return, these drones are charged by the base station and their data off-loaded before the next sortie. Static base stations will have high-speed connectivity to the Cloud, while mobile ones may rely on 4G.

The drones perform on-board real-time Edge computing for navigation [26], combined with control signals from the Fog, which is required for its survival. Additional processing of observations for applications may depend on the spare computing capacity, battery levels and the latency needs of the application. The mobility of the drones and the base allows P2P mechanisms to be effected between the drones using high-speed millimeter wave line-of-sight communication or even between multiple bases. This can be used to move data or processes among them to conserve battery or reduce application latency, such as a mini-drones with constrained resources transmitting data to larger drones while in-flight [22].

IV. FOG COMPUTING PLATFORM ABSTRACTION

Fog computing, even as a concept, is evolving and as a result, there is limited insight on what the *application runtime and management platform* will turn out to be. They have the possibility of incorporating a service model similar to the Cloud as they are not resource-starved like the Edge, but may also inherit some requirements of mobility and performance-sensitivity seen in the Edge [10]. Here, we discuss how the system and application characteristics of Fog computing influence the features and abstractions offered by its platform.

A. Platform Management

One of the key gaps in realizing Fog computing is the lack of a platform ecosystem to design and run applications that can use the Fog, in conjunction with Edge and Cloud layers. There are two key parts to this: a *fabric* for device management, and a *platform* for application runtime management. Clouds expose every resource “as a Service”, and this in part is a reason for its success [10], [21]. Cloud fabrics like OpenStack manage, schedule and instantiate VM images and instances, block and table storage, and networking capacity. Virtualization allows applications and Big Data platforms designed for desktop and cluster computing to work equally well on Clouds.

Device and application management on the Edge is a challenge. Edge computing is done in either *ad hoc*, or tightly coupled to a smart phone platform like Android using sandboxed “apps”. Emerging IoT specifications like CoRE, MQTT and oneM2M enable management, data transfer and control of edge devices. IoT edge management software like Azure IoT Gateway, Amazon AWS Greengrass and VMware Liota are starting to be available. Application platforms like Apache Edgent and MiNiFi try to make logic deployment and migration on the Edge easier. However, these are still evolving.

Fog computing is at an even earlier stage of platform maturity, but can gain from advances in both Cloud and Edge. Slimmed-down versions of Cloud fabrics could manage clusters of Fog devices, though they may have to operate in a wide-area network, while Fogs that are resource-light can make use of platform software for the Edge [8]. An IaaS model for Fog resources would offer the most flexibility, but interoperability of APIs becomes even more important that for Clouds – large-scale Fog deployments will take time, and the ability to use resources from different providers, like cell-phone “roaming”, is essential [15]. Light-weight containers like `lxc` and Docker are likely to be favored for Fog rather than hypervisors, though advances like Ubuntu’s `lxd` may make Linux VMs responsive while also offering live-migration which can be invaluable for mobile Fog resources [3], [10].

B. Application Composition

The execution environment offered by Edge, Fog and Cloud is inherently distributed, and the application space is vast as well. There are many application definition models in use, such as control flows, data flows, and event-driven models [31]. Directed Acyclic Graphs (DAGs) are popular for capturing flow dependencies in complex distributed applications, and are widely used for Cloud and Edge computing [25]. Latency sensitive applications may prefer an event-driven model that react rapidly to changing situations [9].

The ability to encode priorities among tasks and latency guarantees for processing events will be crucial for mission-critical systems. Besides latency, energy and cost are likely to be important application specifications [8], [20]. Geo-fencing policies may also limit the movement of data or services to specific resources.

The data flowing between loosely coupled parts of an application may be based on *streams, micro-batches or files*, again depending on the latency and processing costs. While streams offer a low-latency data transfer for sensor observations, micro-batches have recently grown popular as a balance between low latency and high throughput by amortizing per-tuple overheads³. Files are well suited for exchanging large datasets in batch applications. The ability to define specialized data structures, compression and transport mechanisms for distinct stream types such as video may be necessary too.

State is likely to play a key role in such applications, and the composition models must help include it as a first-class entity. State may be associated with a user, device or session, may need migrate across devices and time, and offer a context for execution [10]. Another unique aspect is the role of *spatial proximity* or location-awareness in determining actions. E.g., proximity of two devices may trigger an action, such as a data transfer, and this may be dictated by their geo-location as well [22]. Some may wish to explicitly couple the application definition with the placement of their constituent tasks, either for performance reasons (e.g., sensing task on edge, processing task on Fog) or for spatial locality to Edge [14].

C. Orchestration and Coordination

The application definition needs to be scheduled and coordinated in order to meet various QoS goals such as latency, energy and monetary constraints. This coordination can be done using different strategies. Three common orchestration models that are relevant in such a multi-layered and distributed resource environment are *centralized, hierarchical* and *peer-to-peer (P2P)*. We also distinguish between *scheduling decisions*, the *flow of control signals* and the *flow of data*, and different coordination models could be applied to these.

Centralized orchestration has a single service, either per application or for the platform, that is located in one of the three resource layers, makes scheduling decisions, and coordinates the transfer of control signals and/or data items. E.g., in [20], a central resource management layer determines the best resources to schedule the incoming tasks by monitoring the Edge, Fog and Cloud resources. This is simple to design but can suffer from high latencies and transfer costs, and is a single point of failure. While this orchestrator often runs in the Cloud (to coordinate across edge devices) or the Edge (to interact with different Cloud services), the Fog layer could offer a sweet-spot for such a centralized coordinator. E.g., when a Fog sits at the gateway between private and public networks, it is in the sole position of being able to judge the behavior of resources in the Edge (private) and the Cloud (public) [32].

A *hierarchical* architecture is a generalization of the centralized model, and allows only vertical communication of data and controls to take place between adjacent layers. This is a

natural fit for Fog computing as it leverages both the bandwidth and latency benefits of the Fog layer in accelerating these flows, as well as the compute benefits closer to the observation source [3], [9], [14], [20], [21], [21]. Often, the Cloud forms the root of this tree and is used for global data aggregation and coordination. Local data analytics is delegated to Cloudlets and further to the edge devices. This allows a federated behavior that has shown to scale.

P2P is a form of distributed coordination that avoids a single point of failure. Here, peers in the same Edge or Fog layer can pass control and data directly among each other [12]. The horizontal communication channels may initially be setup by an entity that has a global picture of the resources. This is typically done at the Cloud or the Fog, or one of the edge devices that serves as a leader. There are simple component based models for composing and executing P2P applications, as well as complex ones that use Distributed Hash Table (DHT) to maintain an overlay network over peers that frequently enter and leave the system. E.g., in [15] the fog devices peer amongst each other for content delivery and service provisioning for improving the performance. Similarly, [18] consider a hierarchical architecture where proximate smart traffic lights coordinate among themselves to send green traffic wave or alert the approaching vehicles.

In a *hybrid* model, there are no strict limitations on the flow of control or data flows, and all layers are seen as having resources of heterogeneous characteristics. While there can be interconnections among resources within each layer (Cloud, Fog, Edge), communication can also take place vertically [10]. This can help in different situations. If a higher layer is not reachable due to network failure or resource unavailability, then the tasks can be accomplished within the lower layers itself though with a graceful degradation of the QoS [24]. Similarly, if a Fog layer is unable to reach an edge device, it can replicate the last state of the application running on it to a different edge device. Such hybrid strategies are likely to require more complex coordination, but can potentially improve the resilience of the application.

D. Managing Dynamism and Mobility

A Fog computing platform has to be responsive to various forms of *dynamism*. As identified before, there can be resource mobility at the Edge and Fog layers. The applications may also impose requirements on logical mobility of processes and data. Further, there may be changes in the data generation rates, network behavior or energy levels of batteries that requires reactive strategies [25]. Thus, the runtime environment should offer monitoring capabilities to determine when such adaptation is required, and provide transparent mechanisms for enacting these changes [20]. This may even require *changing the coordination strategy* from, say, centralized to P2P.

Such *migration* may involve moving just state from one service to another, moving an entire application and its dependencies, or moving a VM or container as a whole. E.g., [10] migrate the user's data from one Cloudlet to another as per the mobility of the user in order to minimize latency. Depending on whether the Edge and or the Fog resources are physically mobile, we may need to preemptively off-load to a different resource or layer when proximity or connectivity is going to be lost, or periodically save state to the Cloud [25].

³Apache Flink, Apache Spark Streaming

For real-time applications or those that have a closed control loop are particularly sensitive to mobility and require careful platform design. Often, these preclude the use of mobile resources, or retain the decision logic in a single resource. E.g., [33], perform video analytics to blur the scenes of movies based on the mental state of the user with a medical condition, but offload the signal processing task to a static Fog server. Batch applications offer the most flexibility and can efficiently make use of mobile Edge and Fog resources to opportunistically offload the tasks or data as and when resources are available. E.g., mini drones may be able to offload their data to larger drones or to a Fog server at the base-station when they come within range [22].

V. RELATED WORK

There have been several early papers that conceptualize the idea of Fog computing and Cloudlets, as has been discussed before. [8] highlight the importance of multi-tenancy on the Fog layer while ensuring mission-critical operation, using a “Foglet” as a software agent to manage the local operations and interface with the Cloud. We go beyond this, and discriminate between edge devices and the Fog layer, and discuss the interaction between all three layers. We also include mobility as first class entity, not just at the Edge but also at the Fog layer. Further, while a centralized, distributed or hierarchical models are mentioned, we contrast between control flows and data flows that may each use a different interaction model.

Cloudlets were introduced in [4], for latency sensitive applications like augmented reality and face recognition applications. They use VMs to deploy and manage applications, either through active migration or of state overlays over existing VM images. They also discuss the business model for Cloudlets, spanning between retain owners and service providers.

Some early research investigates platform and application models for Fog computing. [14] offer a multi-way 3-level tree model where the computation is rooted in the Cloud, resources are elastically acquired in the Cloud and Fog layers, and communication is possible between Cloud and Fog, or Fog and Edge. They consider a strictly hierarchical model. Their example applications do not consider a role for the Cloud either, though their APIs support it. This degenerates to a client-server model between the edges and their Fog parent. Further, the interactions between edge devices and Fog layer should be actively used as well, rather than only vertical interactions. Lastly, while mobility of the edge is discussed, this does not consider when the Fog can be mobile as well.

The role of virtualization in enabling Cloud computing is discussed in [10], and they see a similar role for the Fog as well. They conceive of a VM encapsulating all dependencies for an edge application or user to be hosted on a Cloudlet within 1 hop of the Edge, with this VM moving with the edge user to remain at 1-hop distance. This virtualization architecture should expose API for the developers to offload data and processing, synchronization of data among replicas, discovery of Cloudlet resources, and the migration of the VMs across Cloudlets. However, we take a broader platform view and discuss the possible architectural designs for the Fog.

Others examine specific applications that benefit from the Fog layer. [7] discuss the role of Fog computing for IoT

as Edge computing alone is inadequate handle multiple IoT applications. Fog helps coordinate distributed edge devices and use Cloud resources. They consider Sense-and-actuate and stream processing as two application models, but we argue that more diverse composition and coordination models are needed.

The need of Fog computing for real-time applications such as sensing of gas pipelines, smart agriculture and control systems inside factories has been discussed in [3], based on a hierarchical architecture. But other possible architectural designs and types of applications are missing. Similarly, [21] also discuss the need of Fog computing for smart city applications, and further look into privacy and security issues with Fog computing. What is lacking in these is a broader examination of application characteristics rather than just exemplars that motivate Fog computing, which we address.

A similar effort as our paper has been attempted in [15], but their exercise is limited to mobile devices rather than edge devices at large. They recognize that the Fog can be static or mobile, similar to entertainment systems in vehicles. They note that the Fog can deliver location-aware content unlike Cloud, but this reduces its value to a CDN, only closer to the edge. Moreover, many Cloud services already use geo-location using network or device GPS to offer location sensitive content. Rather, we state that the physical proximity also offers a certain degree of trust by the client of the Fog, and the ability to host rich services, not just content. Their discussion on research problems delves more on the networking and communication aspects than on applications and middleware.

A brief survey of Fog computing concepts has been offered in [9], in which they consider resource poor (edge) devices, and resource rich Cloudlets and Cisco’s IOx. They highlight Augmented reality, Content delivery and Mobile Data Analytics as three motivating applications to reduce the latency delays, and reduce the bandwidth costs. They do not offer any analysis of Fog computing architectural or platform dimensions. They identify network management using Network Virtualization and SDN as enablers of Fog, and QoS, programming APIs, resource provisioning, security and billing as gaps to address. Similarly, [18] present a survey and discuss a hierarchical Fog-based architecture. We discuss several alternative architectural and coordination designs rather than a one-size fits all. They too consider IoT applications but fail to present a taxonomy of the application characteristics that can benefit from Fog as we do.

VI. DISCUSSION AND CONCLUSIONS

A. Challenges

There are several interesting problems that arise in the context of Fog computing, and addressing these can pave way for the technical feasibility.

- *Programmability*: It should be easy for the application developers to run their applications on the Fog. This will require the interface and the programming model to hide the complexities of Fog computing. Fog components should also allow for application-specific customization and optimizations [7], [12].
- *Predicting demand*: One of the major challenges for deployment of Fog devices is to predict the users’ demand

and accordingly determine the quantity of resources required at different locations [15]. Thus it is a challenge to optimally place the fog devices at appropriate locations to serve the requests of nearby mobile users.

- *Power and Network Consumption:* Tasks in Fog devices are distributed across multiple locations compared to more centralized Clouds [7]. Due to the distributed nature of computations, captive local power supply and network connectivity will be important. This can be a challenge for developing countries without 24×7 power supply. Task placement may need to be optimized to minimize the power consumption, or leverage renewables like Solar.
- *Where to push the tasks?* One has to decide how to partition an application and distribute the tasks among the three layers. This depends on the application requirements [9].
- *Security and fault tolerance:* Since Fog resources are geographically distributed and can have multiple service providers, ensuring application and data security is a key challenge [20], [21]. Also, one needs to plan how to handle failures of Fog devices. When the nearby Fog devices are inaccessible, users may need to push their applications to the remote Fog or even the Cloud, which may impact the performance of the applications.
- *Fog providers and billing:* Fog computing services may be provided by Internet, Cloud and Cell-phone service providers, retail merchants, or emerging “gig economies” like taxi or hotel aggregators. Pricing and billing remains a challenge in terms of sustaining a commercial ecosystem or value added services [9]. Moreover, there are possible billing models such as *consumption based* where the users are billed as per their usage or *subscription based* where the users pay a fixed price monthly and can use the Fog network-wide [10].

B. Reality Check

It is still early days for Fog computing [12]. Growth of fog has nowhere near kept up with the growth in edge devices. So the pain-points or gaps of a lack of fog computing have not reached a threshold. These are starting to emerge more in vertically integrated “private” scenarios rather than horizontal, reusable “public” ones, much in the way of Cloud data centers being private to Amazon, Google and Microsoft, before the business model fell into place for commercialization [9].

We do not have large scale deployments of Fog Computing or its commercial operation in a pay-as-you-go on-demand model like public Clouds currently available. But the need for it is growing and we can foresee several existing infrastructure operators evolving to offer Fog services as well. Two prime contenders are *operators of cell-phone towers* who tend to have captive power, communications and space, and can complement this with computing resources [17], and *taxi aggregators* like Uber and Ola who are already offering value added services such as WiFi and grocery delivery, and can incrementally extend this to a Fog server hosted in the boot of their cab. At the same time, emerging infrastructure such as *Smart Cities* will see the organic growth of the Fog layer, initially integrated as part of a vertical domain such as smart power grid or smart transportation (e.g., collocated with a transformer or traffic light, at each city block), and later offered as a horizontal resource shared by multiple city or

commercial services [8], [19], [21]. Such “light-pole” sensing and computing is already starting, as mentioned earlier.

At the same time, not all mobile or IoT applications will necessarily need or even benefit from having a Fog Computing layer. E.g., as Smart Power Meters are rolled out across Los Angeles, the largest public utility in the US, these meters record the electricity usage data at 15 min granularity to support demand-response optimization decisions that are taken periodically [28]. Even with $\approx 4M$ customers, this works out to about 400 GB of data collected each day (assuming a 1 kb payload per observation), or about 4 GB of distributed data collected per 15 min interval. This is meager, even for a megacity like Los Angeles, compared to the 24 GB/sec streamed by Netflix during peak hours ⁴. While there are other Smart Grid applications that may require higher sampling rates and data sizes (e.g., realtime analytics over Phasor Measurement Units (PMU) data [6]), these are still evolving. So not every Smart City or IoT application requires low latency or high bandwidth or complex analytics.

The concept of Utility computing itself existed long before Cloud computing in its current form emerged as a feasible business model, with Grid computing having meanwhile been tried, less successfully, in the academic community [24], [34]. Likewise, sensor networks and pervasive computing concepts existed a decade before the technologies converged to offer IoT at massive scales. Likewise, the reality is that the gestation period for Fog computing may be anywhere from a couple of years to a decade before it matures into a sustainable technology and business model. In fact, one of the early analogies to Fog computing was the *Google Search Appliance*, which offered on-premises search capability for enterprises, with the Cloud only for support. But this died a natural death with network bandwidth improving and the Cloud capabilities like deep-learning far out-stripping what was possible at the GSA ⁵. So the eventual outcome of Fog computing may indeed be different from what one might envision, but offering dimensions for characterization allows this eventual manifestation to fall at some point in the spectrum.

ACKNOWLEDGMENTS

This work was supported in part by grants from the Ministry of Electronics and Information Technology (MeitY), Government of India, and the Robert Bosch Center for Cyber Physical Systems (RBCCPS), IISc.

REFERENCES

- [1] M. Satyanarayanan, P. Simoens, Y. Xiao, P. Pillai, Z. Chen, K. Ha, W. Hu, and B. Amos, “Edge analytics in the internet of things,” *IEEE Pervasive Computing*, vol. 14, no. 2, pp. 24–31, 2015.
- [2] S. Aman, Y. Simmhan, and V. K. Prasanna, “Energy management systems: state of the art and emerging trends,” *IEEE Communications Magazine*, vol. 51, no. 1, pp. 114–119, January 2013.
- [3] M. Yannuzzi, R. Milito, R. Serral-Graci, D. Montero, and M. Nemirovsky, “Key ingredients in an iot recipe: Fog computing, cloud computing, and more fog computing,” in *IEEE CAMAD*, 2014.
- [4] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, “The case for vm-based cloudlets in mobile computing,” *IEEE pervasive Computing*, vol. 8, no. 4, pp. 14–23, 2009.

⁴Evolution of the Netflix Data Pipeline, February 15, 2016, <http://techblog.netflix.com/2016/02/evolution-of-netflix-data-pipeline.html>

⁵http://www.theregister.co.uk/2016/02/10/google_quits_search_appliances/

- [5] S. Aman, M. Frincu, C. Chelmis, M. Noor, Y. Simmhan, and V. K. Prasanna, "Prediction models for dynamic demand response: Requirements, challenges, and insights," in *IEEE SmartGridComm*, 2015.
- [6] J. Interrante and K. S. Aggour, "Applying cluster computing to enable a large-scale smart grid stability monitoring application," in *IEEE HPC-ICISS*, 2012.
- [7] A. V. Dastjerdi and R. Buyya, "Fog computing: Helping the internet of things realize its potential," *Computer*, vol. 49, no. 8, 2016.
- [8] F. Bonomi, R. Milito, P. Natarajan, and J. Zhu, *Fog Computing: A Platform for Internet of Things and Analytics*. Cham: Springer International Publishing, 2014, pp. 169–186.
- [9] S. Yi, C. Li, and Q. Li, "A survey of fog computing: concepts, applications and issues," in *ACM Workshop on Mobile Big Data*, 2015.
- [10] L. F. Bittencourt, M. M. Lopes, I. Petri, and O. F. Rana, "Towards virtual machine migration in fog computing," in *3PGCIC*, 2015.
- [11] R. Ghosh and Y. Simmhan, "Distributed scheduling of event analytics across edge and cloud," *CoRR*, vol. abs/1608.01537, 2016.
- [12] L. M. Vaquero and L. Rodero-Merino, "Finding your way in the fog: Towards a comprehensive definition of fog computing," *ACM SIGCOMM Computer Comm. Review*, vol. 44, no. 5, pp. 27–32, 2014.
- [13] M. Satyanarayanan, G. Lewis, E. Morris, S. Simanta, J. Boleng, and K. Ha, "The role of cloudlets in hostile environments," *IEEE Pervasive Computing*, vol. 12, no. 4, pp. 40–49, 2013.
- [14] K. Hong, D. Lillethun, U. Ramachandran, B. Ottenwalder, and B. Koldchofe, "Mobile fog: A programming model for large-scale applications on the internet of things," in *ACM MCC*, 2013.
- [15] T. H. Luan, L. Gao, Z. Li, Y. Xiang, and L. Sun, "Fog computing: Focusing on mobile users at the edge," *CoRR*, vol. abs/1502.01815, 2015.
- [16] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," *IEEE Pervasive Computing*, vol. 8, no. 4, pp. 14–23, Oct 2009.
- [17] M. Satyanarayanan, R. Schuster, M. Ebling, G. Fettweis, H. Flinck, K. Joshi, and K. Sabnani, "An open ecosystem for mobile-cloud convergence," *IEEE Comm. Magazine*, vol. 53, no. 3, pp. 63–70, 2015.
- [18] I. Stojmenovic, "Fog computing: A cloud to the ground support for smart things and machine-to-machine networks," in *ATNAC*, 2014.
- [19] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *ACM MCC*, 2012.
- [20] A. V. Dastjerdi, H. Gupta, R. N. Calheiros, S. K. Ghosh, and R. Buyya, "Fog computing: Principles, architectures, and applications," *CoRR*, vol. abs/1601.02752, 2016.
- [21] I. Stojmenovic and S. Wen, "The fog computing paradigm: Scenarios and security issues," in *FedCSIS*, 2014.
- [22] S. W. Loke, "The internet of flying-things: Opportunities and challenges with airborne fog computing and mobile cloud in the clouds," *CoRR*, vol. abs/1507.04492, 2015.
- [23] P. Misra, Y. Simmhan, and J. Warrior, "Towards a practical architecture for india centric internet of things," *IEEE IoT Newsletter*, 2015.
- [24] H. Madsen, B. Burtschy, G. Albeanu, and F. Popentiu-Vladicescu, "Reliability in the utility computing era: Towards reliable fog computing," in *IEEE International Conference on Systems, Signals and Image Processing (IWSSIP)*, 2013.
- [25] C. Shi, V. Lakafosis, M. H. Ammar, and E. W. Zegura, "Serendipity: enabling remote computing among intermittently connected mobile devices," in *ACM MobiHoc*. ACM, 2012, pp. 145–154.
- [26] J. Jin, J. Gubbi, T. Luo, and M. Palaniswami, "Network architecture and qos issues in the internet of things for a smart city," in *IEEE International Symposium on Communications and Information Technologies (ISCIT)*, 2012.
- [27] B. Amrutur, et al., "An Open Smart City IoT Test Bed: Street Light Poles as Smart City Spines," in *ACM/IEEE IoTDI*, 2017, (To appear).
- [28] Y. Simmhan, S. Aman, A. Kumbhare, R. Liu, S. Stevens, Q. Zhou, and V. Prasanna, "Cloud-based software platform for data-driven smart grid management," *IEEE/AIP CiSE*, 2013.
- [29] S. M. Amin and B. F. Wollenberg, "Toward a smart grid: power delivery for the 21st century," *IEEE power and energy magazine*, vol. 3, no. 5, pp. 34–41, 2005.
- [30] S. Aman, Y. Simmhan, and V. K. Prasanna, "Holistic measures for evaluating prediction models in smart grids," *IEEE TKDE*, vol. 27, no. 2, 2015.
- [31] Y. Simmhan, R. Barga, C. Van Ingen, E. Lazowska, and A. Szalay, "On building scientific workflow systems for data management in the cloud," in *IEEE e-Science*, 2008.
- [32] M. Aazam and E. N. Huh, "Fog computing and smart gateway based communication for cloud of things," in *International Conference on Future Internet of Things and Cloud*, 2014.
- [33] K. Sadeghi, A. Banerjee, J. Sohankar, and S. Gupta, "Optimization of brain mobile interface applications using iot," in *IEEE International Conference on High Performance Computing*, 2016.
- [34] L. Youseff, M. Butrico, and D. D. Silva, "Toward a unified ontology of cloud computing," in *GCE Workshop*, 2008.