

The Fog-of-Things Paradigm: Road towards On-demand Internet-of-Things

Ruozhou Yu, Guoliang Xue, Vishnu Teja Kilari, Xiang Zhang

Abstract—In this article, we introduce the concept of Fog-of-Things (FoT), a paradigm for on-demand Internet-of-Things (IoT). On-demand IoT is an IoT platform where heterogeneous connected things can be accessed and managed via a uniform platform based on real-time demands. Realizing such a platform faces challenges including heterogeneity, scalability, responsiveness and robustness, due to the large-scale and complex nature of an IoT environment. The FoT paradigm features the incorporation of fog computing power, which empowers not only the IoT applications, but more importantly the scalable and efficient management of the system itself. FoT utilizes a flat-structured virtualization plane and a hierarchical control plane, both of which extend to the network edge and can be reconfigured in real-time, to achieve various design goals. In addition to describing detailed design of the FoT paradigm, we also highlight challenges and opportunities involved in the deployment, management and operation of such an on-demand IoT platform. We hope this article can shed some light on how to build and maintain a practical and extensible control backend to enable large-scale IoTs that empower our connected world.

Keywords—Internet-of-Things, fog computing, Fog-of-Things, on-demand IoT, control plane

I. INTRODUCTION

The Internet-of-Things (IoT) has been recognized as one of the next mega-trends in the technology world. With its ability to interconnect billions of smart things in global scale, IoT is expected to empower innumerable new services and applications that could improve human lives, including smart cities, smart health, smart homes, Industry 4.0, etc. IoT has a huge economic impact: the global IoT market is projected to over 1 trillion dollars in the early 2020s [1].

The vision of IoT is powered by the billions of connected smart things, which virtualize the real-world into digital data that can be transmitted, analyzed, and further utilized. Yet, the massive amount of things has lead to challenges for IoT. On one hand, the current IoT, built upon existing infrastructures such as cellular networks, can hardly accept and process the tremendous amount of data generated by connected things. On the other hand, IoT manageability is also challenged by the huge number of heterogeneous things and the dynamic nature of IoT where things are plugged, unplugged, moving or under failure frequently. It is hard to manage billions of connected things with fine-grained control in real-time. This situation is exacerbated by the stringent requirements of many IoT applications. A majority of IoT applications are *real-time* in nature, meaning that they are designed to analyze continuous

data streams from different locations. These applications have stringent quality-of-service (QoS) requirements, including but not limited to computing power, latency, throughput, loss, robustness, etc.

IoT devices are currently accessed and managed in an ad hoc manner: most deployed things are only visible and accessible to the deployer or owner itself, and have no public access in general. The cause of this is the lack of an on-demand architecture that provides scalable and flexible management as well as uniform and universal access to IoT services. The recently proposed *Cloud of Things* (CoT) paradigm aims to address this issue [2]. CoT provides centralized access and management for smart things in the cloud, therefore achieving a number of benefits including on-demand service, resource pooling, virtualization, etc. However, a cloud-based solution has its limitations. First, it does not resolve the capacity challenge: massive data still needs to be transmitted to the cloud for analysis. Second, due to its long end-to-end latency, the cloud is hard to respond in real-time to frequent network dynamics and failures within vast geographical areas. Cloud-hosted applications also receive unguaranteed QoS in terms of latency, bandwidth and security.

Fog computing is an emerging computing paradigm aiming to address these issues [4], [9]. Different from cloud computing where all computing power is aggregated at a few globally selected locations, fog computing features the deployment of geo-distributed fog nodes across all network hierarchies, and especially in the edge network. These fog nodes provide different levels of capacity and responsiveness to meet various application needs. Fog computing inherits benefits of cloud computing including elasticity and virtualization, while bringing new benefits such as early data resolution, responsive management on the edge, improved latency, robustness and security, etc. However, due to cost and energy consumption issues, fog nodes only have limited capacities, and can only serve things and applications in nearby areas. Collaboration among fog nodes enhances capacity but also complicates system management. Fortunately, cloud computing and fog computing are not incompatible in nature; in fact, they compensate each other's limitations. This has inspired us to seek a unified approach to jointly leverage both cloud and fog computing in IoT.

In this article, we present our understandings on the direction to which the IoT shall evolve. We start from our vision on the emergence of on-demand IoT, describing its necessity and design goals. We then present an on-demand IoT paradigm that jointly utilizes fog and cloud computing, which we call the *Fog-of-Things* (FoT). FoT leverages the distributed and location-aware nature of IoT services, and features a hierarchical and reconfigurable control plane that

Yu, Xue, Kilari and Zhang ({ruozhouy, xue, vkilari, xzhan229}@asu.edu) are all with Arizona State University, Tempe, AZ 85287. This research was supported in part by NSF grants 1421685, 1461886 and 1704092. The information reported here does not reflect the position or the policy of the funding agency.

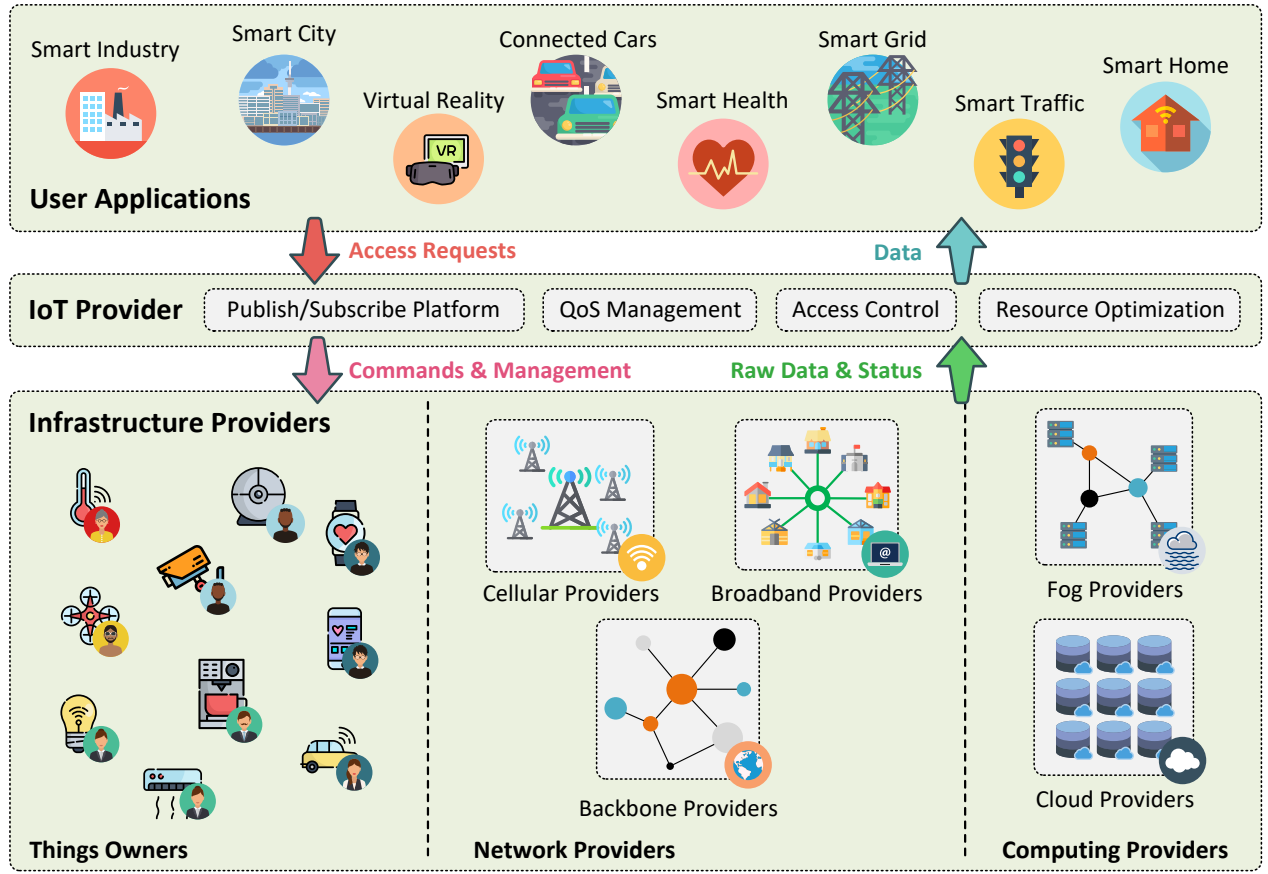


Fig. 1: Three major parties in on-demand IoT: infrastructure providers, IoT provider, and users.

achieves responsiveness, scalability and other design goals. We further describe crucial challenges and opportunities in the deployment, management and operation of FoT. We envision FoT, instead of its ad hoc or cloud-based counterparts, to be one of the pillar stones of our future smart and connected world.

II. ENVISIONING ON-DEMAND IoT

Currently, the IoT infrastructure is mostly deployed and utilized in an ad hoc manner. Various things are produced by different vendors to fulfill similar functions. They are commonly deployed only to power a few applications that belong to the deployer/owner itself, due to lack of proper visibility and accessibility to the public. This has largely buried the true potential of IoT where myriads of distributed things generate massive high-dimensional data that could be used for big data analytics and decision making.

These issues have urged efforts into *on-demand IoT* (also called *service-oriented IoT* [10]), an IoT environment where functionalities (sensing, actuation, data delivery, data analytics, and application hosting) can be provisioned in runtime and remotely accessed by authorized users. It is based on a similar concept as cloud computing, where computing resources are dynamically provisioned and accessed by tenants in an on-demand manner.

Three parties are involved, as shown in Fig. 1. *IoT provider* is who builds and manages the IoT platform, providing various functionalities including virtualization, QoS management,

resource optimization, security, etc. *Infrastructure providers*, such as cloud providers, network providers or things owners, are who participate in the business and provide infrastructure support. Finally, *users*, or *tenants*, are who access and utilize the provided IoT services to develop IoT applications. In general, an on-demand IoT can be built as an overlay upon existing computing and network infrastructures, but can also be built or incremented with self-owned equipment of the providers.

On-demand IoT provides benefits similar to those of cloud computing. From the business perspective, on-demand services could largely reduce capital expenditures (CAPEX) of users by reusing existing infrastructure. A well managed on-demand IoT can reduce operational expenditures (OPEX) as well. For things owners, allowing public access can boost resource utilization, thus increasing their utilities or revenues. For computing or network providers, on-demand IoT enables more flexible pricing options such as pay-as-you-go, which also help increase revenue. From the technical perspective, centralized management helps both infrastructure providers and users. Infrastructure providers can alleviate their overloaded components (network, computing on the edge) by employing smart resource allocation. Users receive guaranteed services via established service-level agreements (SLAs). Finally, a widely accessible IoT platform is also important in inspiring technological innovations, which is beneficial to the entire IoT community.

On-demand IoT is a diverse, large-scale, complex and dynamic environment to build, operate and maintain. IoT is

naturally distributed, and mainly offers location-based services. Hence the same centralized control as in cloud computing is basically not practical. Handling device heterogeneity and dynamicity requires both responsiveness and scalability. In general, system management and optimization is a great challenge both on the infrastructure and on the architectural design. Below, we highlight some design goals of on-demand IoT.

- **Scalability** is probably the most important property of any IoT system. Even a mid-scale IoT needs the ability to manage millions to billions of heterogeneous devices.
- **Virtualization** is crucial in realizing dynamic and elastic services. In general, devices should be exposed only to the minimum extent that their functions can be utilized. Devices with similar functions should be further abstracted using a uniform interface for simple and efficient access.
- **Responsiveness** is more important in IoT environment than in other environments, since a significant portion of IoT applications are time critical. Responsiveness is also crucial in handling system dynamics such as device joining or removal, failure, mobility, etc.
- **Location-awareness** is in the nature of most IoT services. A well-designed IoT system should both provide location-awareness support to applications, and in return utilize it to improve system performance and management.
- **Robustness** is to ensure system functionality during system disturbances such as failures or maintenance. Realizing robustness is specifically crucial in versatile environments like IoT, where disturbances happen frequently.
- **Elasticity** means providing proper scaling and reconfiguration when demands from users change over time. It also means the system can sustain short-term load variations without severe congestion.
- **Security** in IoT is different from it in other environments, mainly because of the constrained nature of IoT devices. Providing native security support to resource-constrained devices is thus an important factor in architectural design. IoT security is vital, since a security breach in IoT can be much more devastating and life-threatening given IoT's ability to monitor and manipulate physical objects.

In the following, we present the design of FoT, an on-demand IoT paradigm. FoT is able to natively achieve several design goals, and also supports realizing the other goals with orthogonal technologies.

III. THE FoT PARADIGM

A. Architecture Overview

Our FoT architecture has four planes, as shown in Fig. 2.

The *data plane* consists of the physical components, including connected things, network switches and routers, and computing nodes. These components perform their functionalities based on upper-plane commands. Due to heterogeneity and dynamicity, the data plane commonly requires frequent reconfiguration and optimization from upper planes.

The *virtualization plane* stands as an intermediary between physical components and decision units in the upper planes. It abstracts heterogeneous physical components into uniform and manageable virtual components.

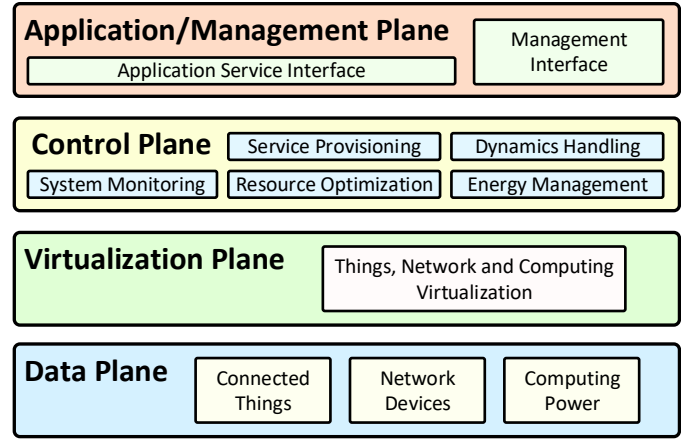


Fig. 2: Four basic planes of FoT.

The *control plane* is the decision core of the architecture. It performs all decision-making tasks in FoT, including component registration, service provisioning, status monitoring and reporting, failure handling, and many more. Our FoT control plane specifically features a recursively built hierarchy of controllers to achieve several design goals of FoT; see Sec. III-D.

The *application/management plane* provides external interfaces of the entire FoT system, and consists of the service interface and the management interface. This plane discloses system services and parameters to authorized users/management teams, and receives application requests and system objectives to be realized by the underlying planes.

B. Data Plane Operation

The data plane consists of the physical components. FoT specifically features the integration of geo-distributed fog nodes, which, in addition to enhancing application performance, also enables local management of other components, as will be detailed in Sec. III-D.

Two key characteristics of the data plane are *heterogeneity* and *dynamicity*. Heterogeneity causes difficulty in automatic management, and can lead to expensive manpower for manual reconfiguration. To address heterogeneity, we add the *virtualization plane* between the conventional control plane and data plane; see Sec. III-C. Dynamicity causes scalability issue and performance fluctuation. We delegate the handling of these dynamics to the control plane to achieve fast and optimized responses; see Sec. III-D.

C. Virtualization Plane Operation

The virtualization plane's goal is to hide the heterogeneity of the data plane. Specifically, for components with similar functions, the system maintains a general *functional template*, which defines the minimum necessary information needed to access and utilize the components. Using the functional template, the system will generate a *functional profile* for each component to describe its function, location, capacity, and other information. E.g., the functional template of a surveillance camera should include its resolution, color profile, location, direction, output format, status messages, command set, etc. These attributes are shared by all surveillance cameras,

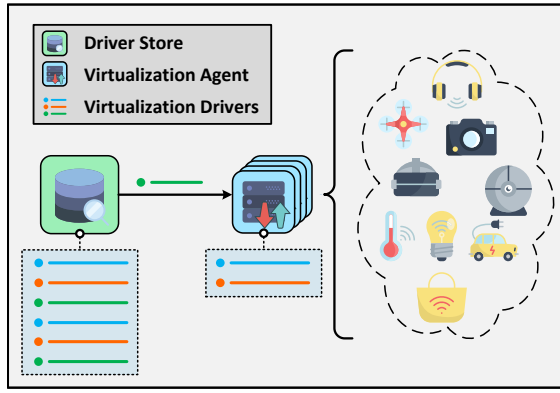


Fig. 3: Virtualization plane design with three main elements: the driver store, virtualization agents, and virtualization drivers.

and hence can be abstracted. Network is commonly abstracted using software-defined networking (SDN). Computing power is commonly abstracted in terms of virtual machines (VM).

Connected things are harder to virtualize than network and computing, as they can be very diverse in functions and specifications. We design the virtualization plane shown in Fig. 3, which consists of three basic elements: the *virtualization drivers*, the *driver store*, and the *virtualization agents*. For all similar components (e.g., the same series of sensors of a vendor), the system maintains a *virtualization driver*, i.e., a module that translates component-specific profile into functional profile of the component. All virtualization drivers are stored in the *driver store*, a centralized database. When the platform introduces a new type of components, e.g., a new sensor model, the corresponding driver is added to the store by the IoT provider. Component virtualization is automatically performed by *virtualization agents*. Each agent keeps a list of locally stored drivers. When a new component is connected, its information is sent to the nearest agent, who will search its local list for the corresponding driver. If the driver is not available, the agent will download it from the central driver store. The agent then performs virtualization for the component as well as subsequent same-type components. These agents are distributed in the network, such as alongside controllers or at the access points. They act as local bridges between the heterogeneous data plane and the uniform control plane.

D. Control Plane Operation

The control plane implements all the management and optimization functionalities. We propose a novel hierarchical control plane, which utilizes the in-network computing power provided by fog computing to resolve the control plane scalability problem.

Hierarchical structure. Our FoT control plane features a hierarchy of controllers that apply control over data plane components in a large geographical area, as shown in Fig. 4. The controllers are organized into a tree structure. At the bottom are leaf controllers, each covering a certain area of connected things and other components (routers, fog nodes, etc.). E.g., a leaf controller can control all components in a smart building or a smart home. On top of that, several adjacent low-level controllers are aggregated and controlled by a parent controller. The parent controller also controls left-over areas between its children's controlled areas. Each controller

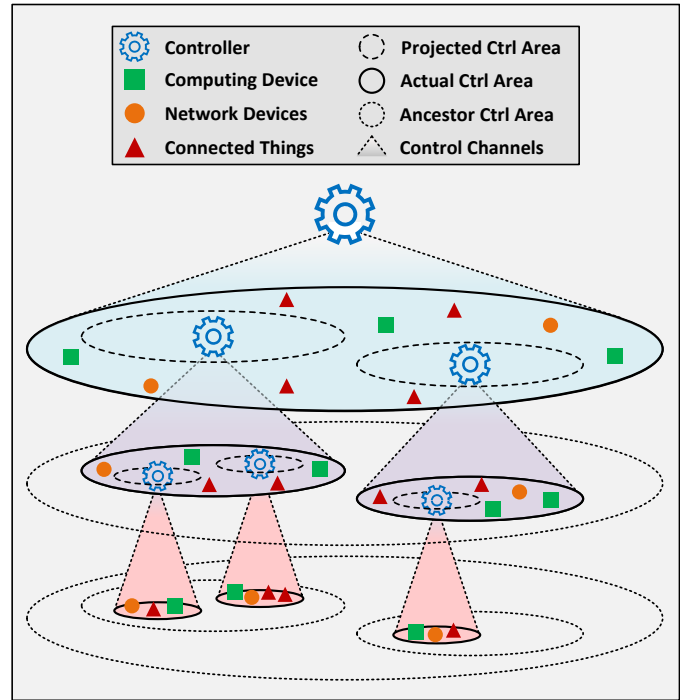


Fig. 4: Our control plane design with hierarchical and recursive controllers. Dashed circles show the projected control area of a child controller of the current level, while dotted circles show the control area of ancestor controllers of current-level controllers. Child controllers are deployed in dense area to alleviate parent load and/or provide better responsiveness.

is located at a computing node within or near the area it controls, which has sufficient computing power to support the controller's operation. The root controller aggregates global information and is commonly located in the cloud.

Recursive operation. Controllers operate in a recursive manner. Each controller (except the leaf controllers) applies both *direct* and *indirect* control to its control domain. Specifically, for components within its child controller's area, the parent controller indirectly issues queries and commands via the child. E.g., if a new device access request is received at the parent controller, the request will be passed to the corresponding child controller, who will process the request and provide the corresponding access to the device if the request is authorized. For components not covered by any child, the controller directly queries and commands the components. The internal logic of our design is to ensure that *control tasks are handled by the lowest-possible level of control*. E.g., routing between devices within a smart building can be directly handled by the leaf controller of the building, without referring to higher-level controllers.

Self-contained reconfiguration. Similar to the recursive architecture for cellular networks [13], the FoT control plane is reconfigurable. Controller assignment is based on the density of components within an area, and can be reconfigured by the parent controller on-the-fly. Furthermore, we argue that controllers should be designed to be *self-contained*. This means that a parent controller can automatically deploy and configure new child controllers at emerging dense areas in its control domain, using its controlled fog computing power, without human intervention. This enables a self-organizing control plane

that can automatically adjust to system load, which is very important in achieving scalability, responsiveness, robustness and elasticity.

Benefits. The benefits of our design is several-fold. First, it achieves scalability by utilizing the location-awareness of IoT services, reducing the states stored at higher-level controllers. E.g., the root controller does not need to store the statuses of most individual devices. Each controller now works on a limited local view of the whole network, which greatly increases the overall scalability of the system. Second, our design improves responsiveness and robustness when facing network dynamics. When a component moves or fails, this event is immediately handled by the direct controller. In case of system-wide optimization, the controller hierarchy can be configured to participate in distributed optimization, which amortizes the control overhead of using a single controller.

E. Application/Management Plane Operation

The application/management plane has two parts: the *service interface* and the *management interface*. The service interface exposes IoT services to tenants, including sensing and actuation by connected things, connectivity, and fog and cloud computing. It accepts application requests from tenants, and translates them into services that will be accommodated by the control plane. Similarly, the management interface exposes system status to the IoT provider. The IoT provider can specify policies through the interface, which will be enforced by the control plane. This enables efficient management without frequently diving into system details.

F. System Function Use Cases

1) *Component Registration:* Components need to be registered to be visible and accessible. Component registration and management should be handled close to the edge to alleviate overhead at higher-level controllers. Initially, the component broadcasts a hello message. The message is sent to the nearest virtualization agent following default network rules. Upon reception, the agent looks up or downloads the virtualization driver, and generates the functional profile of the component. This information is sent to the direct area controller, who creates a *virtual identity* of the component containing its Uniform Resource Identifier (URI) and functional profile. In subsequent operations, the controller will command the component based on its functional profile. Note that the binding among component, agent and controller can be reconfigured in runtime. E.g., if the component moves, a new virtualization agent may take the charge and report to a new controller.

2) *Service Query:* There are two ways that an application can query for a service. First, if the application knows the service URI, a direct request can be issued through the service interface. The request will be broadcast to the entire control plane. If the service is found, all controllers along the broadcast path will jointly establish routing between the servicing component and the application. Second, the application can request for a generic service (e.g., video surveillance) at or near a specific location. In this case, the request will be shipped to the direct controller of the location-of-interest along the control hierarchy, who will then query its local components and find one that fulfills the request. In general, direct queries can be

inefficient due to the need for searching the system. To reduce the overhead of direct queries, the URI can be designed to encode location information. This, however, requires mobility management in the addressing level [6].

3) *Mobility Management and Failure Handling:* FoT handles network events in two steps. First, the event is immediately tackled by local controller for fast response. E.g., the leaf controller will quickly handle a local failure by finding local alternatives to avoid service disruption, such as finding replacement sensors to cover the area of a failed sensor, or alternative routing paths to bypass a router failure. Second, if the event has impact exceeding the local area, or alternatives cannot be found locally, the event will be reported to upper levels for further coordination. E.g., a regional failure may involve multiple controllers to resolve.

Note that our architecture is naturally robust against control plane failures. When a child controller fails, its controlled area is automatically handed-over to the parent, which ensures the continuity of system operation. This way, the only single point of failure is the root controller, which can be backed-up by replicas in the cloud. Still, events are handled by the lowest-possible level of control, which ensures responsiveness and scalability of the system.

IV. CHALLENGES AND OPPORTUNITIES

A. Deployment

An IoT system cannot be built in one night. Incremental deployment enables early utilization of system services, while augmenting the system with new equipment and services during normal operation. Hence the system can be scaled based on provider needs. There are several factors that need to be considered when incrementing the system. First, deploying new devices incurs various costs including deployment, management, energy, etc. Second, investment into different dimensions (new connected things, network devices, or computing nodes) at different locations may result in different improvements of system performance. Therefore, it is advised that the IoT provider optimize its deployment utility based on system-wide measurement of performance.

B. Management

Runtime control plane reconfiguration. One key feature of our design is that controllers can be automatically deployed or revoked based on component density. This ensures elastic control against fluctuating loads and incremental deployment. However, deploying controllers can be costly, especially at dense areas where computing power is already scarce. In this case, the system needs to consider the trade-off between deploying more local controllers to improve system manageability, or devoting more computing power and energy to improving application performance. Finding the optimal deployment and assignment policy subject to capacities and dynamic loads constitutes an optimization problem to be addressed.

Network planning and orchestration. In IoT, network is largely the performance bottleneck due to its limited capacity and long delay. Network planning techniques such as QoS-aware routing [7], traffic engineering [12] and interference

management [15] have each demonstrated their advantages in different network environments. However, applying these techniques in IoT incurs scalability and dynamicity issues. To resolve the scalability issue, proper traffic classification and/or aggregation is needed to reduce control plane states. To resolve dynamicity, both offline and online algorithms need to be developed; the former achieves resource planning in the long run, while the latter provides quick responses to network dynamics. Furthermore, service function chaining should be considered during network planning, which constitutes the network orchestration problem.

Service provisioning and orchestration. Service provisioning is to fulfill application service requests. Services are provisioned in several dimensions: connected thing access, data delivery, and data processing. These can be considered either separately or jointly. E.g., a real-time processing application that analyzes data streams from distributed sensors would require joint consideration of sensor data access, data delivery, (potential) in-network pre-processing, and analytics logic embedding. Factors to be considered include usage costs, network and computing power, energy consumption, QoS (bandwidth, latency, etc.), robustness, elasticity, multi-tenant resource sharing, etc. An important consideration is to utilize geo-distributed fog nodes to host data processing and analytics, to achieve early resolution of the massive data at the edge. A related problem is service orchestration [4], where an application is decomposed into distributed sub-services. Optimization algorithms can be developed, but a general framework that can incorporate different dimensions and constraints is better preferred.

Energy management and optimization. Energy management is a crucial part of IoT [11]. First, a large number of connected things are battery-powered, and hence have very tight energy budget. Second, deployment of connected things tend to be more dense in areas that are already congested with devices, e.g., business districts, urban centers or industrial factories. Energy consumption of the massive things could cause trouble to the energy grid that serves other more critical services, such as lighting or emergency systems. The use of energy harvesters could alleviate the situation, but are not available in scenarios like indoor environments. Proper energy management should jointly consider energy consumption of all components, and utilize various energy sources including power networks, local power generators, energy harvesters, and the smart grid. Both short-term and long-term energy planning is helpful in FoT.

Scalability. With all the flexibility of centralized management comes the concern of scalability. E.g., using SDN as the network controller may suffer from the intrinsic scalability issue of SDN controllers. Our hierarchical control plane serves as a natural remedy to such an issue. The IoT provider can deploy multiple levels of SDN controllers, where each controller controls a local domain, much like the way our FoT controllers work. In fact, our hierarchical control plane can be viewed as a generalization of the existing hierarchical SDN architecture [13], which has already demonstrated the scalability gain of such a design. Nevertheless, scalability will remain a problem in IoT even with such an approach, and surely worths future research and development efforts.

C. Security

Constrained device security. One major challenge in IoT security is the constrained nature of IoT components. Components such as battery-powered sensors or radio frequency identification (RFID) tags have very limited computing power and energy budget [8], and thus are hardly capable of running complex cryptographic algorithms. With the emergence of IoT-related attacks [5], development of constrained security mechanisms is both important and urgent. Yet this field of study is still in its infancy, and requires extensive efforts in the near future.

Infrastructure-assisted security. One way of realizing effective yet inexpensive IoT security is to rely on the platform itself. Such a practice have already been utilized in other environments like data centers or backbones: by deploying security functions within the network, traffic can be checked before reaching the end hosts. Such an approach can be extended to the IoT scenario. E.g., functions that help establish secure channels at access points could greatly alleviate the resource burden on constrained devices, while still receiving most of the benefits of secure transmission. Despite some early efforts [14], there has yet been many researches in this field. We anticipate that infrastructure-assisted security will play a significant role in IoT security.

Privacy. Privacy is of significant concerns in IoT, since a majority of IoT applications are based on locations [3]. While existing location privacy mechanisms address this issue for each single service or application, using a combination of different location-based services could still leak sensitive information. This again urges privacy mechanisms natively embedded into the platform instead of handled by things owners or application developers. IoT can also evade private spaces like homes or factories, which could cause leakage of other sensitive information than location. Proper protection of such information is yet another direction of future research.

V. CONCLUSIONS

In this article, we presented a novel on-demand IoT paradigm named FoT, the *Fog of Things*. The FoT paradigm extends both the data plane and the control plane to the network edge, thus achieving many benefits including scalability, responsiveness, robustness, location-awareness, etc. We designed the FoT architecture which features a flat-structured virtualization plane and a hierarchical and recursive control plane. The virtualization plane achieves universal abstraction of data plane components, while the control plane achieves scalable and fine-grained control utilizing location-awareness of IoT services. We explained the detailed operation of each plane and the system. We also highlighted challenges and opportunities in deployment, management and security of FoT respectively. In general, we envisioned the FoT paradigm to be a major enabler of on-demand IoT, and will play a crucial role in our smart and connected future. Enabling such a future, however, requires extensive future work on the development and implementation of the FoT framework as well as resolving its various issues such as scalability and security.

REFERENCES

- [1] "IoT Market Forecasts," (accessed on May-15-2018). URL: <https://www.postscapes.com/internet-of-things-market-size/>

- [2] A. Botta, W. de Donato, V. Persico, and A. Pescapé, "Integration of Cloud Computing and Internet of Things: A Survey," *Futur. Gener. Comput. Syst.*, vol. 56, pp. 684–700, mar 2016.
- [3] L. Chen, S. Thombre, K. Jarvinen, E. S. Lohan, A. Alen-Savikko, H. Leppakoski, M. Z. H. Bhuiyan, S. Bu-Pasha, G. N. Ferrara, S. Honkala, J. Lindqvist, L. Ruotsalainen, P. Korpisaari, and H. Kuusniemi, "Robustness, Security and Privacy in Location-Based Services for Future IoT: A Survey," *IEEE Access*, vol. 5, pp. 8956–8977, 2017.
- [4] M. Chiang, S. Ha, C.-L. I, F. Rizzo, and T. Zhang, "Clarifying Fog Computing and Networking: 10 Questions and Answers," *IEEE Commun. Mag.*, vol. 55, no. 4, pp. 18–20, apr 2017.
- [5] S. Cobb, "10 Things to Know About the October 21 IoT DDoS Attacks," (accessed on May-15-2018). URL: <https://www.welivesecurity.com/2016/10/24/10-things-know-october-21-iot-ddos-attacks/>
- [6] J. Eriksson, M. Faloutsos, and S. V. Krishnamurthy, "DART: Dynamic Address Routing for Scalable Ad Hoc and Mesh Networks," *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, pp. 119–132, feb 2007.
- [7] J. W. Guck, A. Van Bemt, M. Reisslein, and W. Kellerer, "Unicast QoS Routing Algorithms for SDN: A Comprehensive Survey and Performance Evaluation," *IEEE Commun. Surv. Tutorials*, vol. 20, no. 1, pp. 388–415, 2018.
- [8] H. Hellaoui, M. Koudil, and A. Bouabdallah, "Energy-efficient Mechanisms in Security of the Internet of Things: A Survey," *Comput. Networks*, vol. 127, pp. 173–189, nov 2017.
- [9] P. Hu, S. Dhelim, H. Ning, and T. Qiu, "Survey on Fog Computing: Architecture, Key Technologies, Applications and Open Issues," *J. Netw. Comput. Appl.*, vol. 98, no. September, pp. 27–42, nov 2017.
- [10] V. Issarny, G. Bouloukakakis, N. Georgantas, and B. Billet, "Revisiting Service-Oriented Architecture for the IoT: A Middleware Perspective," in *Proc. ICSOC*, 2016, pp. 3–17.
- [11] N. Kaur and S. K. Sood, "An Energy-Efficient Architecture for the Internet of Things (IoT)," *IEEE Syst. J.*, vol. 11, no. 2, pp. 796–805, jun 2017.
- [12] A. Mendiola, J. Astorga, E. Jacob, and M. Higuero, "A Survey on the Contributions of Software-Defined Networking to Traffic Engineering," *IEEE Commun. Surv. Tutorials*, vol. 19, no. 2, pp. 918–953, 2017.
- [13] M. Moradi, W. Wu, L. E. Li, and Z. M. Mao, "SoftMoW: Recursive and Reconfigurable Cellular WAN Architecture," in *Proc. ACM CoNEXT*, 2014, pp. 377–390.
- [14] S. Nair, S. Abraham, and O. A. Ibrahim, "Security Fusion: A New Security Architecture for Resource-Constrained Environments," in *Proc. USENIX HotSec*, 2011, pp. 2–2.
- [15] M. Noura and R. Nordin, "A Survey on Interference Management for Device-to-Device (D2D) Communication and Its Challenges in 5G Networks," *J. Netw. Comput. Appl.*, vol. 71, pp. 130–150, aug 2016.




Ruozhou Yu (Student Member 2013) received his B.S. degree from Beijing University of Posts and Telecommunications, Beijing, China, in 2013. Currently he is a Ph.D student in the School of Computing, Informatics, and Decision Systems Engineering at Arizona State University. His research interests include network virtualization, software-defined networking, cloud and data center networks, edge computing and internet-of-things, etc.



Guoliang Xue (Member 1996, Senior Member 1999, Fellow, 2011) is a Professor of Computer Science and Engineering at Arizona State University. He received the Ph.D degree (1991) in computer science from the University of Minnesota, Minneapolis, USA. His research interests include survivability, security, and resource allocation issues in networks. He is an Editor of *IEEE Network* and the Area Editor of *IEEE Transactions on Wireless Communications* for the area of Wireless Networking.



Vishnu Teja Kilari (Student Member 2013) received his M.S. degree from Arizona State University, Tempe, Arizona, U.S.A in 2013. Currently he is a Ph.D student in the School of Computing, Informatics, and Decision Systems Engineering at Arizona State University. His research interests include Botnets, Smart Grid security and hardware assisted security.



Xiang Zhang (Student Member 2013) received his B.S. degree from University of Science and Technology of China, Hefei, China, in 2012. Currently he is a Ph.D student in the School of Computing, Informatics, and Decision Systems Engineering at Arizona State University. His research interests include network economics and game theory in crowdsourcing and cognitive radio networks.