

Challenges and Software Architecture for Fog Computing

Zijiang Hao, Ed Novak[†], Shanhe Yi, Qun Li

College of William and Mary, Williamsburg, VA, USA

[†]Franklin and Marshall College, Lancaster, PA, USA

{hebo, syi, liqun}@cs.wm.edu, [†]enovak@fandm.edu

Abstract—In this paper, we give a detailed description of fog computing (also termed edge computing) and show the research challenges and problems. Based on our understanding of these challenges and problems, we propose a software architecture, which is flexible to incorporate different design choices and user-specified policies. Then we report our design of WM-FOG, a computing framework for fog environments that embraces this software architecture, and show the evaluation of our prototype system.

Keywords—System architectures, integration and modeling; Distributed applications

I. INTRODUCTION

The explosive proliferation of ubiquitously connected devices has revolutionized every aspect of human life. However, since most end-user devices, such as smartphones, tablets and wearable devices, have generally weaker CPUs, limited network connectivity, less memory and storage, etc., applications nowadays are usually backed by cloud services. Cloud computing has significantly changed the way we leverage resources for computation, networking, and storage. It provides resources on an on-demand basis, saves expenditures on hardware, and breaks down various technical barriers. Resource pooling in data centers indeed provides more than enough resources for end-user devices. However, moving data from the edge of Internet to the core of Internet, where most data centers are located, is not easy, especially when more and more data nowadays is user-generated content that requires high bandwidth to transmit, according to a recent report [1]. In addition, the unpredictable delay may ruin the user experience of delay-sensitive applications, such as human-computer interfaces, emergency services, and real-time games. While we believe that cloud computing will still be a mainstream computing paradigm in the future, the rapid development of Internet and pervasive mobile devices has called for a new computing paradigm that can overcome the inherent drawbacks of cloud computing, such as unpredictable latency, bandwidth bottlenecks, lack of mobility support and location awareness, and so on. To this end, fog computing (also termed edge

computing) has been proposed as a non-trivial extension of cloud computing. By providing elastic resources at the edge of network, fog computing can better support a variety of emerging applications.

Fog computing has been defined from several perspectives [2], [3], and similar concepts such as cloudlets [4], mobile-edge computing [5] and mobile-cloud computing [6] have also been proposed. We have given a more general definition of fog computing in our previous work [7] as “*a geographically distributed computing architecture, with a resource pool that consists of one or more ubiquitously connected heterogeneous devices (including edge devices) at the edge of network, and not exclusively seamlessly backed by cloud services, to collaboratively provide elastic computation, storage and communication (and many other new services and tasks) in virtualization isolated environments to a large scale of clients in proximity.*”

Fog computing will benefit several relevant domains, including mobile/wearable computing [8], Internet of Things (IoT) and big data analytics, in reducing latency, increasing throughput, consolidating resources, saving energy, and enhancing security and privacy [9]. In IoT, fog computing can provide unified interfaces and flexible resources to accomplish heterogeneous computational and storage requests. In virtual reality (VR), a 3D VR gaming headset has to be cable-connected to a high-end server for low latency on processing complex 3D graphics. Fog computing can fulfill the low latency need for VR users in proximity, and can save expenditures on extra hardware. In big data analytics, huge volumes of data are generated at the edge of network. Fog computing supports edge analytics, which can reduce the delay of data analytics and decrease the cost of data transmission and storage.

In this paper, we introduce fog computing, a new computing paradigm that extends cloud computing. Fog computing promises performance benefits such as low latency and quick response time in various application scenarios. We compare fog computing and cloud computing in detail, and list a number of research

challenges and problems. Based on our understanding of these challenges and problems, we propose a software architecture, which is flexible to incorporate different design choices and user-specified policies, and highlight WM-FOG, a computing framework for fog environments that embraces this software architecture. We also conduct experiments on our prototype system to show that WM-FOG can work effectively and efficiently in real-world fog environments.

II. COMPARISONS BETWEEN FOG AND CLOUD

There are several key differences between fog computing and cloud computing. Cloud servers are usually rack-mounted high-end servers located in large warehouse-like data centers. Centralized cloud servers allow for replication, load balancing, failure recovery, power management, and easy access to failed hardware for repairing and replacement. For this reason, the reliability of cloud services can be held at a very high standard. The exact opposite can be expected in fog computing. Fog nodes are geographically distributed, scattered all over the edges of Internet, and logically decentralized in that they are maintained by different organizations. Consequently, fog nodes are not as reliable as cloud servers, and physically locating a failed fog node and repairing it is more difficult and costly. Many financial and time costs, such as those related to power and system configuration, cannot be amortized as they would be with cloud computing. Another key insight is that the network connectivity to fog nodes cannot be guaranteed. An unreachable fog node cannot fulfill any request even if its computational hardware is fully functional. Simple tasks like regular testing and auditing of hardware are immensely more complicated and costly in fog computing, due to necessary coordination among different organizations, geographical distribution, and unreliable network connectivity.

Scheduling tasks in fog computing is complex compared with that in cloud computing. A fog computing application is typically spread over a) the client's mobile device, b) one of potentially many fog nodes, and occasionally c) a back-end cloud server. Therefore, deciding where to schedule computational tasks in fog computing is more difficult. For cloud computing applications, the latency is usually predictable. For fog computing applications, however, deciding which fog node to use alters the latency that the user will experience. Beside the unpredictable round-trip time, slow hardware and low bandwidth also affect the user-perceived latency. Meanwhile, some tasks, such as aggregate caching, may benefit from running in the back-end cloud. Another problem is that it is unclear where the scheduling should be done. Entrusting the client devices to perform the

scheduling opens the possibility of malicious users abusing the system. Fog nodes may act selfishly or may not always be aware of tasks running on the client devices. The back-end cloud may introduce unnecessary latency to the scheduling program. In short, in fog computing, more factors must be considered in deciding where and when to schedule tasks to provide the best user experience.

Fog nodes are maintained independently by many different organizations, which is in sharp contrast to the back-end cloud owned and maintained by a single organization. This means that fog nodes cannot be trusted as easily as cloud servers. Users can more easily trust cloud computing because the organizations that provide cloud services are well-motivated to invest in resilient security and privacy measures. Fog computing, on the other hand, is implemented by many independent agents. These various owners may not maintain the same rigorous privacy and security standards, let alone the high standard fulfilled by cloud computing. As such, users will have a much more difficult time trusting different fog nodes.

Fog nodes are heterogeneous, where there is no guarantee that the nodes will contain similar resources. In fact, quite the opposite can be expected; fog nodes, owned and maintained by different organizations, usually have vastly different RAM capacity, CPU performance, storage space, and network bandwidth. This is in sharp contrast to cloud computing, in which it is common for one organization to own all of the cloud servers. To ease the burden of application deployment, hardware management and resource sharing, cloud servers usually exhibit much less heterogeneity. Furthermore, fog nodes are smaller and less powerful than large cloud servers. While a cloud server may be one of many high-end, powerful rack-mounted servers, fog nodes are usually deployed in small batches using desktop-class machines, re-purposed computing appliances such as routers and gaming consoles, and small collections of rack-mounted servers. Because of the heterogeneity and generally weaker hardware in fog computing, many of the differences we have outlined so far are exacerbated.

Fog computing aims to establish a new tier of mobile computing, in which constraints on energy and hardware resources can be relaxed by nearby fog nodes. Users can effortlessly offload computation to nearby fog nodes, and can transparently and seamlessly move computation from one fog node to another. Mobility is a key feature of the fog computing paradigm, and applications deployed on fog infrastructure need to always take this into account. This differs from cloud computing, in which applications are deployed in only one cloud at a time, unless the need for scaling is beyond the capacity of a single cloud provider. A situation that is rarely seen in

cloud computing but may be common in fog computing is that users may connect to the network only briefly while moving. Consider the scenario in which fog nodes are placed along roadways and users temporarily connect to them to acquire traffic and weather conditions ahead. Another example is that users move from one fog node to another when they are leaving their office and heading to a different building on a college or corporate campus for a meeting.

III. RESEARCH CHALLENGES AND PROBLEMS IN FOG COMPUTING

In this section, we discuss the challenges and problems in fog computing. We have identified papers on the concept of fog computing with a joint effort of our previous work [10], [11], [7], [12] and other existing work [13], [14], [15]. We refer the reader to these references for an integral view of the state-of-the-art on fog computing.

Fog computing is a novel computing paradigm, which demands a new programming model. We need to design intuitive and effective tools and frameworks for developers, helping them orchestrate dynamic, hierarchical and heterogeneous resources to build compatible applications on diverse platforms. To take task scheduling and migration for example, various research questions may arise. How can we provide a simple abstraction for developers to mark tasks that can be migrated? What choices and preferences should be left to users? How can we allow developers to specify migration rules on various devices? Furthermore, we should avoid forcing developers to reimplement functionalities that will likely be common, such as distributed caching, workload balancing, system monitoring, and so on.

Fog computing introduces a variety of new and interesting scheduling challenges. Since tasks can now be moved between different physical devices (i.e., client devices, fog nodes, and back-end cloud servers), scheduling is much more complex. Some of the research questions are listed as follows. For fog nodes with heterogeneous hardware, is it acceptable to trade energy for reduced latency? Should a process running on a fog node be interrupted when the user moves toward another fog node? How should tasks be scheduled considering a) latency, b) energy consumption, c) mobility, and d) existing workload? Where should the scheduling program be executed? What are the benefits of jointly scheduling tasks? Beside these research questions, some other concerns must also be taken into account. For example, security and privacy considerations are complex in fog computing, and tasks from sensitive applications should be scheduled on more trustworthy nodes. Furthermore, on traditional desktop and server machines, completely fair scheduling (CFS) dominates the landscape. For fog

computing, however, this algorithm may not be ideal, as different fog nodes may have different hardware resources, and some tasks (e.g., those making the user wait) may be more important than others (e.g., background services such as the backup/snapshot functionality). What other scheduling algorithms can be used to optimize the factors that are important for highly mobile computing, such as low latency?

Data management for fog computing applications also introduces new challenges. Perhaps the ideal abstraction for both users and developers is a “global storage,” which can always be accessed, has an infinite size, and yet performs with the speed of information stored locally. With fog computing, this dream may finally be realized. However, how to implement such a storage system is still an open question. What efficient algorithms can be used to shuffle data among devices? How can prefetching be best implemented to achieve the lowest latency? What namespace scheme should be used? How can sensitive and encrypted data be cached privately and effectively? Furthermore, energy consumption and network usage must be conserved on mobile devices, as they typically have energy limits enforced by limited battery technology and data limits enforced by mobile carriers.

An attractive aspect of cloud computing is the automatic discovery of services. As fog nodes differ from location to location, users can arrive at a new location and take advantage of the various services provided by the fog nodes in that particular location. Given that this feature relies on the prudent deployment accomplished by service developers, implementing service discovery protocols in fog computing can be quite challenging. Moreover, service provisioning is usually done dynamically in fog computing, i.e., new virtual machines (VMs) are orchestrated on the spot when a particular service is needed. This raises many research questions. When should services be started and stopped? What is the best way to balance workload? Should VM-based or container-based virtualization be used? It may even be possible to predict what services will be needed and provision them in advance, before users even arrive. What methodologies should be used to efficiently provision services for hundreds, thousands, and millions of users? Some services are deployed on fog nodes to aggregate information from nearby client devices. What is the best way to split workload for these services, with regard to the energy efficiency on the client side?

In cloud computing, data consistency can be achieved by coordinating the cloud servers in the data centers where the cloud is deployed. In fog computing, however, things become complicated. When writing data objects in a fog environment, it is necessary to not only coordinate the back-end cloud servers, but also invalidate the cached data on the fog nodes as well as on the client devices

if strong data consistency is needed. This may result in deteriorated write performance, which weakens the benefits of using fog nodes as the write cache servers. On the other hand, fog computing also provides opportunities of achieving data consistency more efficiently than cloud computing. For example, if the write requests on a data object are sent to only one fog node during a certain time period, which we envision is a common case in fog computing, the system may temporarily transfer the ownership of the data object from the cloud to the fog node. By doing this, data consistency can be achieved on the fog node, which promises better write performance than cloud computing, as the fog node resides at the edge of network. Nevertheless, to fully exploit the opportunities of achieving data consistency in fog computing is still challenging and requires plenty of research efforts.

Finally, and perhaps most importantly, fog computing and IoT have significant privacy and security concerns. Due to the heterogeneous nature of both, security and privacy are usually cast aside in order to achieve general functionality and interoperability. In other words, encryption and strict privacy policies make it more difficult for arbitrary devices to exchange data. Therefore, many manufacturers today simply do away with these features. Moreover, encryption algorithms and security protocols, which are notoriously complex, are often implemented or configured with mistakes, leaving sensitive user data exposed to attackers. This problem is further exacerbated by the dispersed ownership of fog nodes. Fog nodes are usually owned by different parties, such as universities, corporations, commonwealth organizations, and personal households. Some fog nodes may even be jointly owned by two or more parties. Users approaching a fog node may be weary of the services provided by these parties, due to their vastly differing motivations. As we have seen with online social networking, collection and resale of private user data is highly valuable to corporations. Meanwhile, these parties also need authentication protocols to protect themselves against Sybil accounts, distributed denial of service (DDoS) attacks, and other malicious activities. It will be important in the future to make fog computing applications preserve user privacy, provide rigorous security guarantees, and address the needs of all the parties involved.

Exploring solutions to the aforementioned problems is critical in realizing the many benefits promised by fog computing. It is our goal in this work to expose these design choices in detail, so that developers can more easily figure them out in their implementations. This is a first step towards identifying reasonable solutions to these problems.

IV. WM-FOG OVERVIEW

Based on our understanding of the aforementioned challenges and problems, we propose WM-FOG, a computing framework for fog environments. The design of WM-FOG embraces a software architecture, which is flexible to incorporate different design choices and user-specified policies. More specifically, WM-FOG provides a flexible way to define *workflows* that can be easily deployed and executed on fog-based systems. By properly scheduling the workflows on the system entities (i.e., client devices, fog nodes, and back-end cloud servers), WM-FOG can take advantage of the fog computing paradigm and achieve considerable performance enhancement. Furthermore, and most importantly, WM-FOG provides a way to customize policies on the workflows, through which developers can help the system make even better use of the underlying hardware resources.

A. Workflow Examples

To define a workflow, the developer needs to specify its data and computation. We call the data *data items*, and the computation *transitions* in WM-FOG. Each workflow contains one or more data items and zero or more transitions.

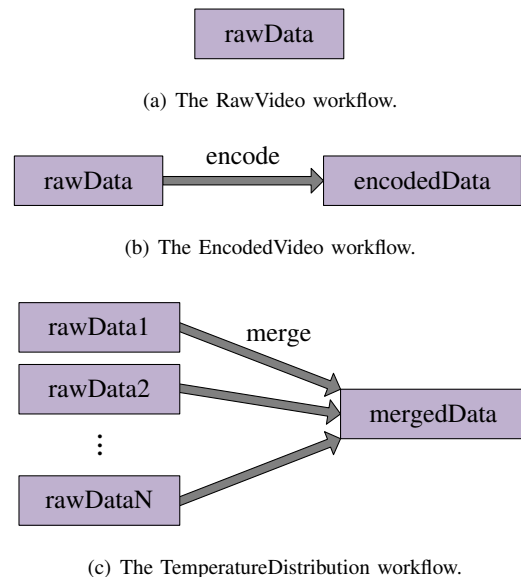


Fig. 1. Workflow examples.

Figure 1(a) illustrates a simple workflow, which is called RawVideo. This workflow contains only one data item, depicted as “rawData” in the figure, and no transition. The only data item rawData represents the raw video data that the WM-FOG system has received from a client device.

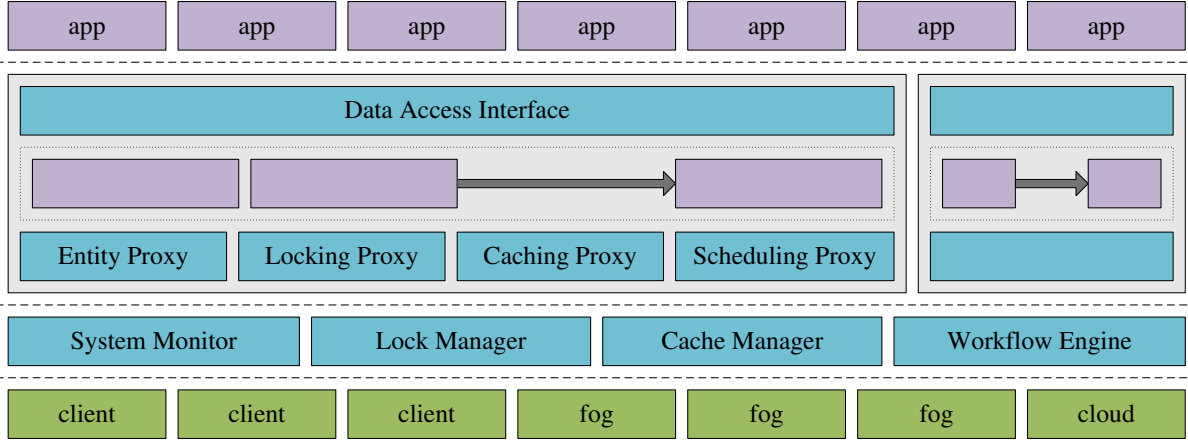


Fig. 2. WM-FOG software stack.

Figure 1(b) illustrates a slightly more complicated workflow, which is called *EncodedVideo*. This workflow contains two data items, *rawData* and *encodedData*, and one transition, *encode*. The *encode* transition takes *rawData* as input and generates *encodedData* as output. Clearly, in this workflow, the WM-FOG system receives raw video data from a client device, encodes it, and stores the encoded data for future use.

Figure 1(c) illustrates a workflow involving multiple writers, which is called *TemperatureDistribution*. Suppose there are N fog nodes covering different regions, and each fog node has a number of temperature sensors deployed in the region that it covers. Each temperature sensor continuously uploads the ambient temperature data to the fog node it belongs to, and the fog node in turn forwards the data to the back-end cloud. The cloud receives data from the N fog nodes, merges them, and stores the merged data for future use.

B. System Architecture

Figure 2 depicts the architecture of WM-FOG. There are four layers in the figure. The top layer is called the *application layer*, where user applications reside. User applications initiate workflow instances by writing input data to them, and receive results by reading output data from them. The next layer is called the *workflow layer*, where workflow instances reside. Each workflow instance exposes a *Data Access Interface* to user applications, through which its data items can be accessed. Moreover, each workflow instance has four proxies, i.e., the *Entity Proxy*, the *Locking Proxy*, the *Caching Proxy*, and the *Scheduling Proxy*. These proxies can be used to implement user-specified policies on workflows. Under the workflow layer is the *system layer*, where the system components, i.e., the *System Monitor*, the *Lock Manager*, the *Cache Manager*, and the *Workflow Engine*, reside.

These system components implement the fundamental mechanisms of WM-FOG, and workflow instances can communicate with them through the proxies in order to apply user-specified policies. The bottom layer is called the *entity layer*, as the system entities (client devices, fog nodes, and the cloud) reside in this layer.

C. Customizing Workflow Policies

WM-FOG provides a workflow-defining language for developers. More specifically, developers can specify the data items and transitions for each workflow they are defining through this language. Furthermore, they can selectively implement the *callback functions* of the data items and transitions to define their own policies. To define a policy on a workflow, the developer is supposed to invoke the workflow's proxies in the callback functions, informing the system components of her suggestions on how to handle the workflow under various conditions. Note that the developer can only provide her suggestions, but not control the behavior of the system components.

1) *Implementing Synchronization Policies:* A developer can implement her own synchronization policy on a data item, by programming its callback functions. In these callback functions, the developer is supposed to invoke the *Caching Proxy* of the workflow to communicate with the *Cache Manager*, providing her suggestions on how to synchronize the data item.

For example, suppose the developer who has defined the *RawVideo* workflow shown in Figure 1(a) has implemented a synchronization policy on the *rawData* data item, which eagerly synchronizes the first 20 MB of data to the back-end cloud, while keeps the remaining part of data on the fog node and synchronizes it lazily. By doing this, user applications that read the data item can get served immediately. Meanwhile, the *Cache Manager* will spontaneously synchronize the remaining part of

data when the system has spare hardware resources, guaranteeing that the data item can be fully synchronized as soon as possible. This synchronization policy has the same effect on serving read requests as the default synchronization policy, which eagerly synchronizes the whole data item to the back-end cloud, but imposes less burden on the system, which is critical when the system is handling a burst of workflow requests.

2) *Implementing Locking Policies*: A developer can implement her own locking policy on a data item, by programming its callback functions. In these callback functions, the developer is supposed to invoke the Locking Proxy of the workflow to communicate with the Lock Manager, providing her suggestions on how to manage the locks on the data item.

For example, the *rawDataI* data items ($I = 1, 2, \dots, N$) shown in Figure 1(c) are written by multiple writers (i.e., temperature sensors). Suppose the *rawData1* data item can be written by only one writer at any time, which requires a locking mechanism to coordinate the write operations upon it. A simple way of implementing such a locking mechanism is to maintain a write lock for the data item in the back-end cloud. Before a writer writes the data item, it has to acquire the write lock from the cloud, while after the data item has been written, the writer needs to return the write lock to the cloud. Using the cloud as the centralized lock server is necessary when different writers try to acquire the same write lock from different fog nodes. However, as previously described, the *rawData1* data item will only be written by temperature sensors belonging to the same fog node. In such a case, using the cloud as the centralized lock server will impose unnecessary overhead on the write operations. The developer can invoke the Lock Proxy in the data item's callback functions, informing the Lock Manager that she suggests the write lock be maintained on the fog node rather than in the cloud. By doing this, the Lock Manager will try to maintain the write lock on the fog node, which can improve the write performance on the data item in most cases. Note again, however, that the developer cannot really control the behavior of the Lock Manager, i.e., if the Lock Manager considers that the write lock should be maintained in the cloud, it will do so, rather than unconditionally follows the developer's suggestion.

3) *Implementing Migration Policies*: When defining a transition, the developer needs to specify its input data items as well as their *trigger thresholds*. For example, the encode transition shown in Figure 1(b) has only one input data item, *rawData*. Suppose the developer has specified that the trigger threshold of the *rawData* data item is 1024 KB when defining the encode transition. In such a case, the Workflow Engine will automatically invoke the *onTrigger()* callback function of the encode

transition whenever the *rawData* data item has enqueued 1024 KB of data. The *onTrigger()* callback function is the place where the developer implements the transition's main logic.

The execution of the *onTrigger()* callback function is atomic in WM-FOG. In other words, the Workflow Engine may migrate a transition between two consecutive executions of the *onTrigger()* callback function, but will never do so during the execution of it. Data that needs to be transferred in a migration should be defined as member variables of the transition. The developer should also provide the getter and setter functions for these member variables.

A developer can implement her own migration policy on a transition, by programming its callback functions excluding *onTrigger()*. In these callback functions, the developer is supposed to invoke the Scheduling Proxy of the workflow to communicate with the Workflow Engine, providing her suggestions on when and how to migrate the transition.

For example, the encode transition shown in Figure 1(b) should be triggered mainly on fog nodes. This is because the WM-FOG system can leverage the computational power of fog nodes to achieve better performance, given that the encode transition has a good compression ratio. Nevertheless, it may be unreasonable to always trigger the transition on fog nodes, especially when they are fully loaded. For this reason, the developer may implement a migration policy on the encode transition, informing the Workflow Engine that she suggests migrating the transition to the back-end cloud if the fog node cannot execute it in 0.5 seconds. By doing this, some workload on fully loaded fog nodes can be offloaded to the back-end cloud, and the overall system performance can thus be improved.

V. EVALUATION

In this section, we give some preliminary results to show that WM-FOG can leverage the fog computing paradigm to enhance the system performance when handling workflow tasks (i.e., workflow instances).

A. Testbed Setup

We build a testbed for our experiments. The testbed consists of five servers, one of which is more powerful than the others. The more powerful server is used as the back-end cloud server, while the others are used as fog nodes. The cloud server has an 8-core Intel i7 CPU with a clock speed of 4.00 GHz and 16 GB of main memory. Each fog node has a 4-core CPU with a clock speed of 2.83 Hz and 4 GB of main memory, and is directly connected to the cloud server through a 1000 Mbps network link. To simulate a real-world fog

environment, we set the upper bound of the network bandwidth between each fog node and the cloud server to 40 Mbps, and the latency to 10 ms (i.e., the round trip time is 20 ms), according to the results reported by [7]. We also deploy our first-step implementation of WM-FOG on this testbed.

B. Benefits of Using Fog

We first evaluate to what extent fog computing can help when handling WM-FOG workflow tasks. To this end, we simulate a scenario in which RawVideo workflow tasks shown in Figure 1(a) are handled by our system. Each RawVideo task has a total data size of 200 MB, and is sent from the client device to the fog node at a transmission rate of 8 Mbps. On each fog node, the arrival intervals of RawVideo tasks follow an $N(10 \text{ sec}, 4 \text{ sec}^2)$ distribution. We cache the 200 MB rawData of each RawVideo task on the fog node, while synchronize only the first n MB of data to the cloud. The value of n is varied from 0 to 200 in our experiments.

Figure 3 illustrates the latency results, the throughput results, and the network usage results of these experiments. From these results, we can see that a smaller synchronization size (i.e., n MB) produces shorter latency, higher throughput, and lower network usage. Despite the fact that the synchronization size cannot be too small for providing seamless data accessing services, these results demonstrate the benefits of using fog computing in WM-FOG, as tuning the synchronization size is only possible in fog environments.

C. WM-FOG Performance

WM-FOG makes decisions on how to handle workflow tasks based on 1) the default policies, and 2) suggestions from developers. In other words, the default policies are part of the fundamental mechanisms of WM-FOG, which are supposed to provide graceful performance for fog environments.

To evaluate how well our prototype system works, we conduct the following experiments. We use only one fog node and the cloud server in these experiments. Once again, we simulate the scenario in which 200 MB RawVideo tasks are handled by our system. A user-specified policy that at least the first 20 MB of rawData should be eagerly synchronized to the cloud is applied to the RawVideo tasks. In the first experiment, we disable the default synchronization policy, so that only the user-specified policy is enforced. In the second experiment, we enable the default synchronization policy, so that the System Monitor monitors the network usage of the fog node. If the System Monitor detects that there is spare network resource between the fog node and the cloud, it will inform the Cache Manager, which will in turn try to synchronize more data for the RawVideo tasks.

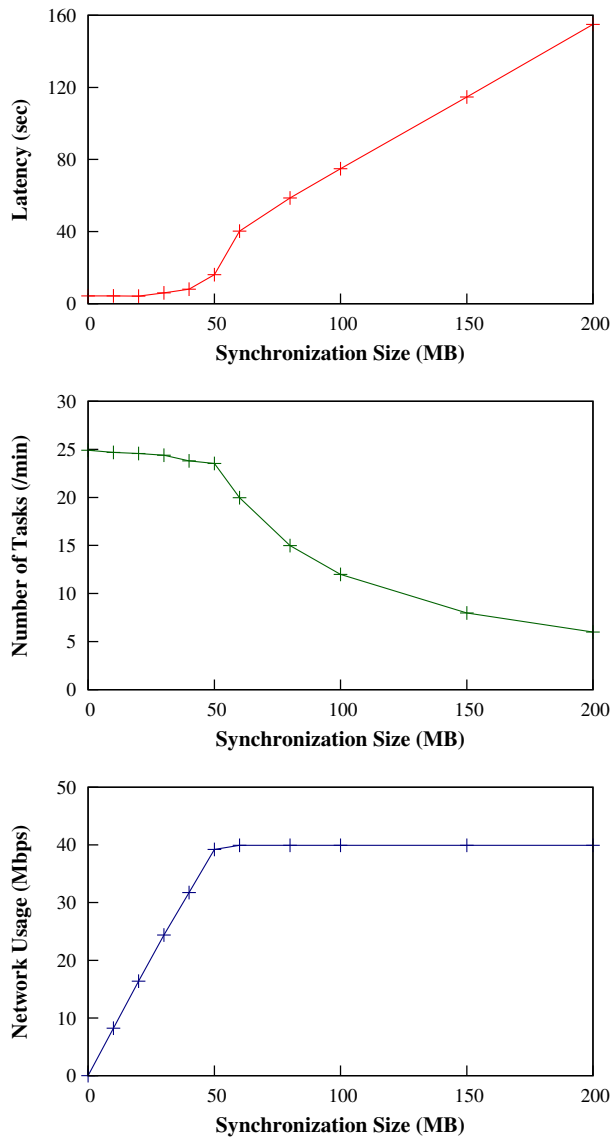


Fig. 3. Performance measurements of using fog.

Figure 4 illustrates the synchronization size results and the network usage results of these experiments. Clearly, when the default synchronization policy is enabled, the system can make better use of the network resource, and the burden of fully synchronizing the RawVideo tasks in the future can thus be reduced. These results demonstrate that WM-FOG is an efficient computing framework for fog environments.

VI. CONCLUSION

In this paper, we introduce fog computing, a new computing paradigm that extends cloud computing. We compare fog computing and cloud computing in detail, and list a number of research challenges and problems in fog computing. Based on our understanding of these

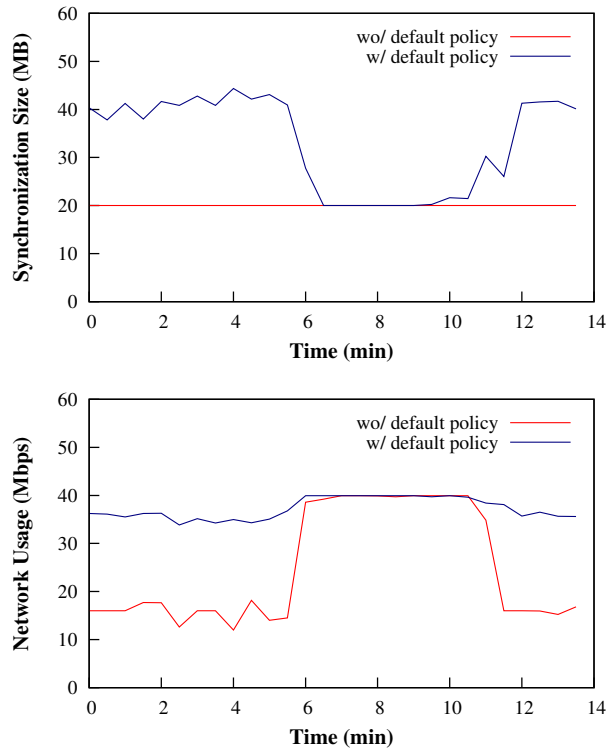


Fig. 4. Performance measurements of WM-FOG.

challenges and problems, we propose a software architecture, which is flexible to incorporate different design choices and user-specified policies. Then we report the design of WM-FOG, a computing framework for fog environments that embraces this software architecture. Evaluation on our prototype system demonstrates that WM-FOG can work effectively and efficiently in fog environments.

REFERENCES

- [1] Sandvine, "Global internet phenomena report," 2015.
- [2] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *Proceedings of MCC*, 2012, pp. 13–16.
- [3] L. M. Vaquero and L. Roderer-Merino, "Finding your way in the fog: Towards a comprehensive definition of fog computing," *SIGCOMM Computer Communication Review*, vol. 44, no. 5, pp. 27–32, 2014.
- [4] M. Satyanarayanan, Z. Chen, K. Ha, W. Hu, W. Richter, and P. Pillai, "Cloudlets: at the leading edge of mobile-cloud convergence," in *Proceedings of MobiCASE*, 2014, pp. 1–9.
- [5] ETSI, "Mobile-edge computing," 2014.
- [6] N. Fernando, S. W. Loke, and W. Rahayu, "Mobile cloud computing: A survey," *Future Generation Computer Systems*, vol. 29, no. 1, pp. 84–106, 2013.
- [7] S. Yi, Z. Hao, Z. Qin, and Q. Li, "Fog computing: Platform and applications," in *Proceedings of HotWeb*, 2015, pp. 73–78.
- [8] Y. Cao, S. Chen, and D. Brown, "Fast: A fog computing assisted distributed analytics system to monitor fall for stroke mitigation," in *Proceedings of NAS*, 2015, pp. 2–11.
- [9] T. Zhang, "Fog boosts capabilities to add more things securely to the internet," <http://blogs.cisco.com/innovation/fog-boosts-capabilities-to-add-more-things-securely-to-the-internet>, 2016.
- [10] S. Yi, C. Li, and Q. Li, "A survey of fog computing: Concepts, applications and issues," in *Proceedings of MoBiData*, 2015, pp. 37–42.
- [11] S. Yi, Z. Qin, and Q. Li, "Security and privacy issues of fog computing: A survey," in *Proceedings of WASA*, 2015, pp. 685–695.
- [12] Z. Hao and Q. Li, "Edgestore: Integrating edge computing into cloud-based storage systems," in *Proceedings of SEC*, 2016.
- [13] I. Stojmenovic, "Fog computing: A cloud to the ground support for smart things and machine-to-machine networks," in *Proceedings of ATNAC*, 2014, pp. 117–122.
- [14] W. Shi and S. Dustdar, "The promise of edge computing," *IEEE Computer Magazine*, vol. 29, no. 5, pp. 78–81, 2016.
- [15] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet of Things Journal*, 2016.

Zijiang Hao is currently working toward the Ph.D. degree in computer science under the supervision of Prof. Qun Li at the College of William and Mary. His research interests include mobile-cloud computing, fog/edge computing, geo-distributed storage systems, and consensus algorithms.

Ed Novak is an assistant professor of computer science at Franklin and Marshall College. His research interests include cybersecurity and privacy on smart mobile devices. He received his Ph.D. degree in computer science from the College of William and Mary in 2016.

Shanhe Yi is currently working toward the Ph.D. degree in computer science under the supervision of Prof. Qun Li at the College of William and Mary. His research interests include mobile/wearable computing and fog/edge computing, with the emphasis on the usability, security and privacy of applications and systems.

Qun Li is a professor of the Department of Computer Science at the College of William and Mary. He holds a Ph.D. degree in computer science from Dartmouth College. His research interests include wireless networks, sensor networks, RFID, pervasive computing systems, and fog/edge computing. He received the NSF Career Award in 2008.