

Fog Function: Serverless Fog Computing for Data Intensive IoT Services

Bin Cheng, Jonathan Fuerst, Gurkan Solmaz
NEC Laboratories Europe, Heidelberg, Germany

Takuya Sanada
NEC Solution Innovators, Ltd., Tokyo, Japan

Abstract—Fog computing can support IoT services with fast response time and low bandwidth usage by moving computation from the cloud to edge devices. However, existing fog computing frameworks have limited *flexibility* to support dynamic service composition with a data-oriented approach. Function-as-a-Service (FaaS) is a promising programming model for fog computing to enhance flexibility, but the current event- or topic-based design of function triggering and the separation of data management and function execution result in *inefficiency* for data-intensive IoT services. To achieve both flexibility and efficiency, we propose a *data-centric programming model* called *Fog Function* and also introduce its underlying orchestration mechanism that leverages three types of contexts: data context, system context, and usage context. Moreover, we showcase a concrete use case for smart parking where Fog Function allows service developers to easily model their service logic with reduced learning efforts compared to a static service topology. Our performance evaluation results show that the Fog Function can be scaled to hundreds of fog nodes. Fog Function can improve system efficiency by saving 95% of the internal data traffic over cloud function and it can reduce service latency by 30% over edge function.

Index Terms—serverless computing, fog computing, edge computing, IoT services, data-centric programming model, function-as-a-service

I. INTRODUCTION

With the proliferation of IoT devices such as sensors, cars, drones, and robots, these *IoT devices* not only produce lots of data but increasingly consume the output of machine learning powered data processing pipelines to take actions in a timely fashion [1]. Usually, data producers and consumers are linked via *IoT services* that implement the data processing logic to transform raw data into actionable results. Previously, central clouds have been used as the main underlying infrastructure for hosting such IoT services. However, due to the requirements of short response time and low bandwidth use of many services associated with connected vehicles, drones and cameras, there is a strong need to move data processing from the cloud to the network edges that are close to both producers and consumers [2]. This paradigm shift is generally labelled fog or edge computing, which involves computing in both cloud and edge environments [3]. For consistency we will use the term *fog computing* throughout the paper.

Various fog computing frameworks exist, such as Azure IoT Edge [4], AWS Greengrass [5], EdgeX [6], and Baidu OpenEdge [7]. However, their programmability for IoT services is limited in terms of flexibility due to the below reasons.

- 1) The design and deployment of services is *bound to specific edge devices*. This is an edge-oriented approach and requires service developers to statically define which service module should be deployed on which type of edge device. The set of service modules running on an edge device is fixed after deployment. However, for situation-aware IoT applications, different service modules need to be triggered dynamically at the network edge according to the availability and mobility of IoT devices. This requires a *data-oriented approach*.
- 2) Existing fog computing frameworks have *poor support for data-intensive IoT services*. For example, most existing fog computing frameworks use a topic-based pub/sub interface, such as MQTT, for communication between different edge service modules and require a manual configuration of the data routing path. This is problematic when a running task instance at the edge needs to be migrated to the cloud or another edge because of device mobility, workload fluctuation or when we need to add or remove service modules dynamically due to changing business needs. Therefore, *IoT services require a dynamic service composition*.

Recently, *serverless computing* [8] is promoted by major cloud providers to support Function-as-a-Service (FaaS) computing in a lightweight, dynamic and event-driven manner. At first glance, this makes it a good fit for the dynamic characteristics of fog computing. However, serverless computing frameworks only deal with the execution management of functions, completely separating it from data management. This separation benefits the simplicity of serverless computing, but has drawbacks for data-intensive batch and stream processing applications [9], [10]. For fog computing, where data locality is paramount, this separation is a deal breaker and results in poor system efficiency.

In this paper, we design and implement serverless fog computing to support data-centric IoT services in an edge-cloud environment. We address the limitations of existing serverless computing and fog computing frameworks, improving their efficiency and flexibility. Specifically, we propose a fog function programming model and a context-driven orchestration runtime system to enable serverless fog computing. Our main technical contributions are listed as follows:

- We design *Fog Function* as an enhanced Function-as-a-Service programming model, which relaxes the lifetime and resource constraints of traditional cloud functions and allows

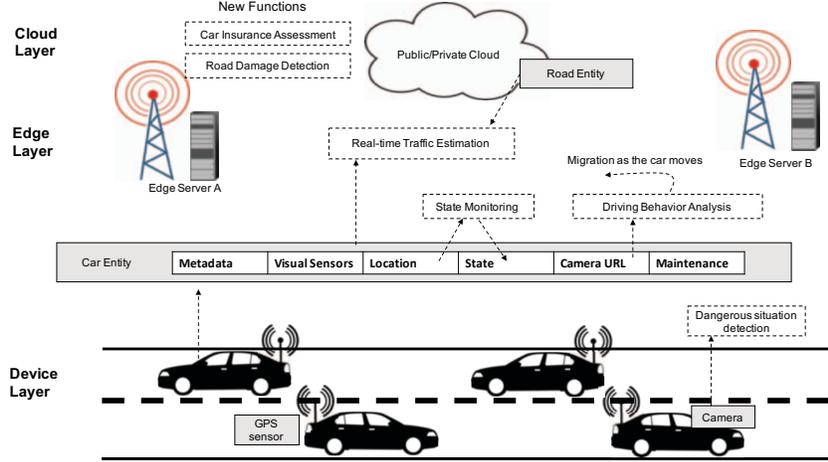


Fig. 1: The connected car use case.

the seamless move from code to data or from data to code. Based on the validation and analysis of two use cases, we show that Fog Function is more flexible and efficient than traditional event-based cloud functions to support serverless fog computing for data-intensive IoT services.

- We propose a context-driven orchestration mechanism to dynamically and automatically trigger, configure, and optimize the deployment of Fog Function in cloud and edge environments. Our approach not only supports mobility aware task migration and but also achieves better load balance across edge nodes as compared to edge-oriented approaches applied in current fog computing frameworks.
- We introduce the detailed mechanism of task deployment and migration to support context-driven Fog Function and report the performance results in terms of latency and scalability.

II. MOTIVATION AND GAP ANALYSIS

A. Motivating Use Case

Several domains, such as smart cities, automotive, and smart manufacturing benefit from serverless fog computing. For example, with a growing number of connected cars, there is an emerging demand to deploy IoT services to enhance driving experience and safety by leveraging the data produced by connected cars and other data sources. Figure 1 depicts such scenario with four IoT services based on data from four entity types: *Road*, *Car*, *Camera*, and *GPS sensor*:

- (S1) *State monitoring*: This service takes the location updates of a car and then checks whether the car is moving in a normal state; in the end it updates the state of the car entity.
- (S2) *Driving behavior analysis*: If the car is in an unusual state, this function will be triggered for a detailed inspection of the driver’s behavior, based on the captured image from a camera in the car.
- (S3) *Real-time traffic estimation*: Estimate the real-time traffic information aggregated at various levels, e.g., per road, per district, per city.

- (S4) *Dangerous situation detection*: Detect any dangerous situation on the road and then update the road entity in order to inform the other drivers behind on the same road.

B. Diversity and Dynamics of IoT Data and Workload

In the car use case the diversity and dynamics of IoT data and workload are reflected by the following observations. Assume that all data are represented as entities while the workloads of IoT services are represented by data processing tasks. Each task takes some entities in, performs some internal data processing, and then produces some outputs to create new entities or update existing entities.

- (O1) *Small vs. big entity*: Entities processed by IoT services differ in size. For example, a car entity can contain lots of information about the car, such as location, manufacture information, embedded sensors, cameras, maintenance information, and other metadata. A road entity might be small.
- (O2) *Static vs. dynamic entity*: Some entities are static or do not change frequently, while others change frequently. E.g., the location of a car entity changes constantly.
- (O3) *Small vs. big task*: The required computation per task differs. For example, the state monitoring task is much more light-weight than the image-based driving behavior analysis.
- (O4) *Short vs. long task*: The lifetime of tasks differs and is often bound to the availability of input data. For example, the driving behavior analysis is only triggered if the car entity state is set to be “abnormal” by the state monitoring task.
- (O5) *Normal vs. urgent task*: Different tasks might have different priorities. For example, S4 has higher priority than all the other services and it needs to have more resource or even exclusive resource usage at computing edges since the available resource at each edge is usually limited.
- (O6) *Existing vs. new task*: New services might be added on the fly at runtime. For example, road damage detection can be deployed to assess the road condition and then trigger timely road inspection and maintenance; suspect detection can be launched by the law enforcement office to search and track

suspicious attackers during an emergency situation. These new tasks need to reuse raw data published by devices and intermediate results generated by other existing tasks.

C. Gap Analysis

To realise these IoT services the following two expectations need to be met: First, in the design phase, service providers want to have the flexibility to add, remove, update, and compose services on the fly as their business evolves over time. Second, during the operating phase, these services should be able to run seamlessly and efficiently across geo-distributed clouds and edges managed by infrastructure providers. To meet both expectations, we identify the following gaps that a fog computing framework needs to address:

(G1) *Data discovery and routing: from topic-based to content-based*: Raw data and intermediate results should be forwarded to different tasks based on their needs. Existing fog computing frameworks use a topic-based pub/sub interface, such as MQTT, to configure the data routing paths between different tasks. However, representing all entity data with topics is inefficient because some task might only need part of the whole entity data (O1). Thus, to efficiently discover and forward any required data to a task, the management and discovery of IoT data needs to be content-based.

(G2) *Function triggering: from per event to per selected entities*: In existing serverless computing frameworks, functions are invoked per event with limited execution time and memory size. This is not suitable for data-intensive IoT services (O3 and O4). First, it is difficult for service designers to know the execution time and memory size required by a function during the design phase. Second, the input of a function can be a single entity update or a stream of updates (O2) and lifetime is usually associated with the availability of input data. Triggering a function with the availability of its inputs not only avoids the effort to explicitly define the triggering event, but also allows service developers to follow data-centric design principles.

(G3) *Function execution: from data→code or code→data to code↔data*: Existing serverless computing frameworks separate data management from the function execution environment, always moving data into the execution environment for function execution (data→code pattern). On the other hand, existing fog computing frameworks such as Azure IoT Edge move cloud functions to the data located at the edges. This follows a code→data pattern. With regards to our observations O2, O3, and O4, the workload of various IoT data processing tasks is highly diverse and changes over time, requiring a dynamic and transparent placement of data and code.

(G4) *Function composition: from event-oriented or edge-oriented to data-centric*: Observation O6 indicates a strong need for function composition. In existing serverless computing frameworks such as OpenWhisk [11], service developers need to customize a series of event triggers and rules to link multiple functions together. Existing fog computing frameworks allow service developers to link functions at each edge by manually configuring the topic-based data routing

path between them. However, for data-intensive IoT services, a data-centric approach is more efficient and flexible because it can directly take advantage of the data dependency between different functions and perform function composition with a global view of the entire data layer, rather than a subview of each edge.

III. FOG FUNCTION FOR SERVERLESS FOG COMPUTING

To fill the aforementioned gaps, we propose a new data-centric function programming model called *Fog Function*. We also introduce the underlying context-driven service orchestration mechanism of the programming model.

A. High Level System View

Figure 2 shows the high level view of our system for the orchestration of fog functions. The system consists of a number of *fog nodes*, each of which runs a *Broker* and a *Worker*. A management node runs two centralized components, namely *Discovery* and *Orchestrator*. Each node is a Virtual Machine (VM) or physical host deployed either in the cloud or at edges. All fog nodes form a hierarchical overlay based on their configured GeoHash IDs. All data in the system are represented as entities saved by a Broker and indexed by the centralized Discovery for discovery. The data can be raw data published by IoT devices, intermediate results generated by some running data processing tasks, or available resource data reported by fog nodes. When a fog function is registered, Orchestrator will subscribe to the input data of the fog function to Discovery. Once the subscribed data appear or disappear in the system, Orchestrator will be informed and then take orchestration actions accordingly, which will be carried out by an assigned worker.

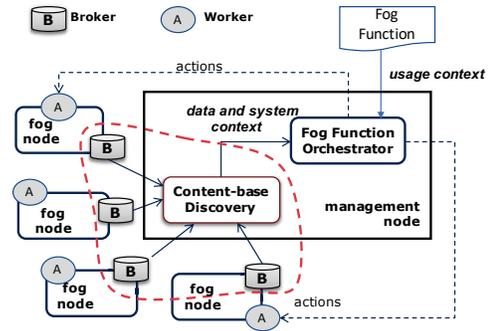


Fig. 2: High level system view

Notice that with our approach the orchestration is also driven by events but limited to only two pre-defined events: entity “appear” or “disappear”. It also provides a declarative interface for service designers to easily annotate which data should be used to trigger which function with a customizable granularity and some other high level orchestration intentions. Therefore, service designers can fully focus on the data aspect without explicitly defining the triggering events. Figure 3 shows the relationship between the identified gaps and our

proposed technology. More details are provided in the following sections.

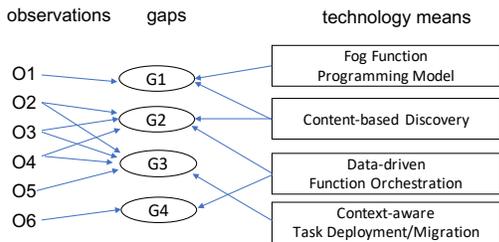


Fig. 3: The mapping between gaps and observations/proposed means

B. Fog Function Programming Model

Each fog function is presented as an entity as well, but it is annotated with the following attributes.

- *Name*: a unique identity of the function.
- *Operator*: the name of a data processing operator. The operator is implemented as a dockerized application based on the interface of fog function described below. The specified operator is instantiated at runtime by a Worker as a task with its configured inputs and outputs. The task is deployed in a dedicated Docker container running on a fog node.
- *Inputs*: a set of selected inputs required by the operator to do internal data processing.
- *Outputs*: the entity type generated by the operator.
- *Geoscope*: the geoscope to be applied when selecting input data for this fog function.
- *Priority*: the priority of this fog function, which will be taken into account by workers to decide how to assign their limited resources to different functions at fog nodes.
- *SLO*: the expected Service Level Objective, which is defined as various optimization goals, for example, minimizing the latency to produce outputs, maximizing the accuracy of generated results, or minimizing the bandwidth usage across fog nodes. Different SLO leads to different task deployment plans.

The last four attributes are optional. The inputs are the key for Orchestrator to decide when to trigger the fog function and how to create its tasks. Each input is further defined with the following information.

- *SelectedType*: the entity type of this selected input.
- *AttributeSet*: the required attributes of the selected entity.
- *Constraints*: the filters to further select input entities based on some specific attribute values.
- *GroupBy*: the granularity to control how many tasks should be instantiated and how the selected input entities should be assigned across its tasks. It can be defined as “per entityID”, “per entityType”, or “per attributeValue”.
- *Scoped*: this could be true or false and it is used to decide whether geoscope should be applied to select input data when the geoscope is defined with the fog function.

The designed interface for developers to program an operator is shown as below. The first parameter *entity* is the received data for internal processing. The other parameters, publish, query, and subscribe, are the callback functions for the internal function code to interact with the data management layer via a nearby broker assigned by Discovery, such as, publish generated entity data, query or subscribe additional information.

```
function(entity, publish, query, subscribe)
```

As compared to the existing FaaS programming models like *Cloud Function* in the existing cloud-based serverless computing frameworks or *Edge Function* in the existing edge-centric fog computing frameworks, Fog Function has a similar interface for developers to write function code, but it provides some unique annotations for the underlying runtime system to efficiently trigger, execute, and compose functions over cloud and edges in a flexible and transparent manner. The comparison results are summarized in Table I from different perspectives.

C. Content-based Discovery

The context management layer consists of a network of Broker(s) and a centralized Discovery. It is designed to provide a global view for all system components and running tasks to query, subscribe, and update context entities via the unified and standardized data model and communication protocol, namely NGSI [12]. It plays a very important role to support the orchestration of Fog Function. As illustrated by Figure 4, in our design a large number of distributed Brokers work in parallel under the coordination of the centralized Discovery. The Discovery component can be used to discover both devices, intermediate data, and available resources at fog nodes as well.

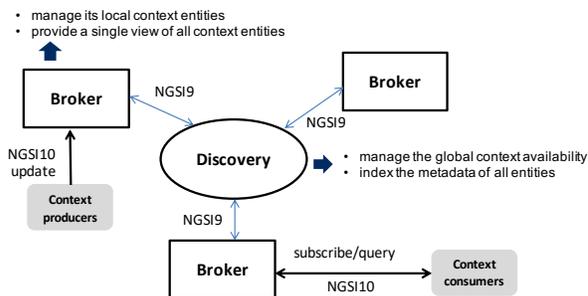


Fig. 4: Context discovery with two-layer design

As compared to the topic-based data management in existing systems like MQTT-based Mosquitto or Apache Kafka, our two-layer context management design has the following features: 1) separating context entity data and their availability; 2) providing separated and standardized interfaces to manage both context data (via NGSI10 [12]) and context availability (via NGSI9 [12]); 3) supporting not only ID-based and topic-based query and subscription but also geoscope-based or attribute-based query and subscription.

TABLE I: Comparison with existing function-based programming models

Differentiator	Function-based programming models		
	<i>Cloud Function</i>	<i>Edge Function</i>	<i>Fog Function</i>
Execution environment	centralized cloud	each edge	cloud and edges
Input and output bindings	one event	per topic	selected entities
Configuration	none	none	tunable parameters
Task granularity	none	none	definable
Trigger	per event	per edge	availability of selected entities
Execution pattern	data \rightarrow code	code \rightarrow data	code \leftrightarrow data
Migration	yes	no	yes
Priority	none	none	yes
Service level objective	none	none	definable

D. Data-driven Function Orchestration

Figure 5 shows the major procedure for Orchestrator to orchestrate fog functions based on the update notification of context availability of their input data, provided by Content-based Discovery. More specifically, the following four basic orchestration actions are designed to dynamically orchestrate tasks for each registered fog function.

- **ADD_TASK**: To launch a new task with the given configuration that includes the initial setting of its input streams. When launching a new task, the Worker first fetches the Docker image for this task and then launches and configures this task within a dedicated Docker container. After that, the Worker subscribes the input entity to the context management system on behalf of the running task so that the input streams can be received by the running task; in the end, the newly created task is reported back to the orchestrator.
- **REMOVE_TASK**: To terminate an existing running task with the given task ID. When terminating an existing task, the Worker not only stops and removes its corresponding Docker container, but also unsubscribes its input streams so that the context management system does not end up with lots of unavailable subscribers.
- **ADD_INPUT**: To subscribe to a new input stream on behalf of a running task so that the new input stream can flow into the running task.
- **REMOVE_INPUT**: To unsubscribe from some existing input stream on behalf of a running task so that the task stops receiving entity updates from this input stream.

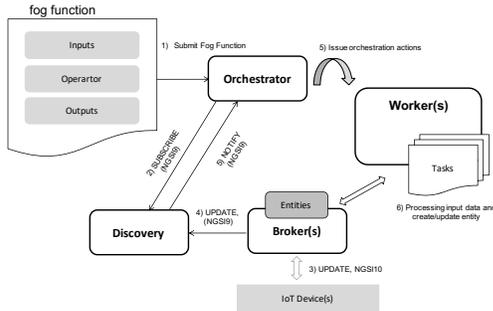


Fig. 5: Data-driven orchestration

E. Context-aware Task Deployment/Migration

Overall the orchestration of Fog Function leverages the following three types of context information.

Data context: the structure and registered metadata of available data, including both raw sensor data and intermediate data. Based on the standardized and unified data model and communication interface, namely NGSI, our system is able to see the content of all data generated by sensors and data processing tasks in the system, such as data type, attributes, registered metadata, relations, and geo-locations.

System context: available resources at each fog node. The resources in a cloud-edge environment are geo-distributed and they are dynamically changing over time. As compared to cloud computing, resources in such a cloud-edge environment are more heterogeneous and dynamic.

Usage context: high level usage intentions defined by service designers to indicate what their fog functions should be used in the system, such as which type of results is expected under which type of QoS within which geo-scope.

Task migration is the combination of removing an existing task on Worker A and adding a new task on another Worker B. Currently, we only support the migration of stateless tasks, meaning that the tasks do not hold any persistent internal state and terminating or restarting them does not lead to any faulty state or result. More specifically, our design can allow seamless task migration across cloud and edges in the following three cases.

- **Cloud \rightarrow Edge**: for example, if a new edge node joins the system and it is close to one edge device, a task running in the cloud can be migrated from the central cloud to this new edge node in order to save bandwidth.
- **Edge \rightarrow Cloud**: for example, when an edge node becomes overloaded since it has to handle the workload from lots of devices in the same region, it can start to migrate some existing tasks to the cloud in order to keep enough resource for the other urgent tasks.
- **Edge \rightarrow Edge**: for example, when a mobile device such as a connected car moves from one region (R1) to another region (R2), it might find another nearby edge node that can provide lower latency. In this case, it is better to migrate the task to the edge node in R2. In this case, tasks are migrated in order to adapt to the movement of the mobile devices.

IV. IMPLEMENTATION AND USE CASE VALIDATION

Fog Function has been applied into the open source fog computing framework FogFlow as a new programming model to enhance its programmability. Originally, FogFlow can orchestrate dynamic data flows over cloud and edges using a service topology [13]. The service topology statically defines the logical data processing flow of an IoT service and is triggered on demand by the requests from the consumer side, but it does not support the composition of multiple service topologies and it is not flexible to handle use case requirements which may change over time. Unlike service topology, Fog Function is simple and flexible and it is triggered when its input data becomes available. FogFlow can automatically chain different functions and allows more than one Fog Function to handle new data items. In the end, the entire execution graph can be automatically triggered, composed, and managed as data arrives. From the design perspective, Fog Function is more flexible than the service topology, because the overall processing logic of an IoT service can be easily changed over time by adding or removing functions when the service processing logic needs to be modified for new use case requirements.

Using Fog Function, we are able to easily realize a *smart parking* use case, which was difficult to achieve with the service topology programming model. This use case is implemented together with our European project partner, University of Murcia, based on the real scenario of Murcia City. In Murcia, there are two types of parking sites, regulated parking zones that are operated by the city government and can provide historical information of how parking slots are used per day, and private parking sites that are operated by private companies and can provide real-time availability of parking spots. By utilizing these two types of data sources and other public transportation information, our smart parking service can provide real-time and personalized parking recommendations for drivers.

As illustrated in Fig. 6, we just need to design and implement dedicated fog functions for each physical object involved in the use case. For example, one fog function for each public site to predict how many parking spots are available per 10 minutes based on their historical information; two fog functions for each connected car, one to estimate its arrival time according to the traffic situation on the way and the other to calculate at which park site the driver can get a parking spot on arrival. The deployment of those fog function instances are on the edge node close to their input data sources so that FogFlow can reduce more than 50% bandwidth consumption and also provide real-time parking recommendation for each driver.

V. PERFORMANCE EVALUATION

Our performance evaluation has been conducted by using a set of virtual machines from Google Cloud. The whole system setting includes 3 parts: 1) a client machine that can simulate a set of IoT devices; 2) a set of fog nodes, of which each is a

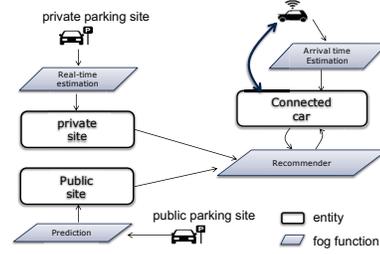


Fig. 6: Design of fog functions for a smart parking use case

standard VM with 2 vCPU and 7.5 GB memory (“n1-standard-2”); 3) the cloud part that includes just one more powerful VM with 8 vCPUs and 7.2 GB memory (“n1-highcpu-8”), running the two centralized components, namely Orchestrator and Discovery. To trigger the designed data processing tasks defined by service topology or Fog Function, we simulate a set of IoT devices for each test case. For each test case, we carry out 10 runs of tests. In each test, we start a number of simulated devices and keep them running for 10 minutes.

A. Latency

Figure 7(a) shows the results of startup latency measured in the following three scenarios: 1) without launching the actual task (named as “task-not-launched”); 2) the docker image of the dummy task is not fetched in advanced (named as “fetch-image-and-launch-task”); 3) the docker image of the dummy task is already fetched (named as “only-launch-task”). From the measurement result, we can see that the big part of the startup latency happens with fetching the docker image from public docker registry. Launching a docker container also requires about 2 seconds. These two parts of latency are also related to the size of the required docker image. The bigger the docker image is, the longer it will take. However, the actual time taken by Orchestrator to make its orchestration decisions is very short, less than 100 ms. This result indicates that, by leveraging multiple Workers to launch tasks in parallel, we can scale up the system easily, even just using a centralized Orchestrator.

We also measure the latency of migrating one task from one Worker to another Worker. By design, migrating a task is done by two orchestration actions: terminating the existing task and starting a new task. For stateless tasks, these two actions can be carried out in parallel. Figure 7(b) shows the latency results of terminating/starting/migrating a task. We can see that the latency of migrating a task is about 2 ~ 2.5 second, nearly equal to the latency of starting a task. This is because the two orchestration actions (starting a task and terminating a task) are performed in parallel and starting a task is much slower than terminating a task.

B. Scalability

We also measured how the FogFlow system can be scaled up by adding more fog nodes in terms of throughput. The throughput is defined as the average number of launched tasks

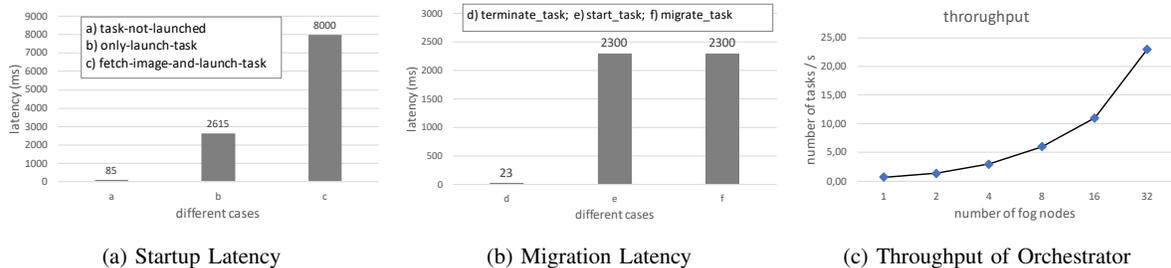


Fig. 7: Initial measurement results of the FogFlow system

per second. Figure 7(c) shows the result. In the test, we assume that the docker image is already fetched in advanced. From the result, we can see that the throughput increases linearly with the number of fog nodes. With 10 fog nodes, FogFlow can handle ~ 8 tasks per second. As seen in Figure 7(c), the main bottleneck is actually the underlying docker engine, which can only launch 1~3 docker containers per second. A single Orchestrator can handle the orchestration decisions for at least 32 FogFlow Workers. In the future, we will utilize the unikernel-based visualization technology, such as Unikraft [14], to avoid this bottleneck. In addition, to further scale up the system, we need to decentralize the Orchestrator and Discovery as well.

C. Efficiency

We evaluate system efficiency for the Fog Function based approach in terms of cross-node traffic and service latency and compare it with the other two approaches: Cloud Function and Edge Function. Cloud Function represents the existing severless computing approach for the cloud environment. With this approach function is triggered by event based on topic and the data management environment is separated from the execution environment. Edge Function represents the existing fog computing frameworks that program services per edge based on topic. We change the system setting of FogFlow and the specification of Fog Function to simulate both Cloud Function and Edge Function. For topic-based pub/sub, we use type-based and subscribe to the entire entity; for the separation of data management and function execution, we deploy brokers with Discovery and workers with Orchestrator within two groups of VMs. We simulate 1000 connected car entities with two types of entity lengths (126 bytes for the small size case and 1682 bytes for the big size case) to trigger a simple speed estimation function. Each car reports its current location every second and each run of test is 10 minutes. The cross-node traffic is the total amount data transferred across VMs and the service latency is the delay from when the raw data are published to when the result is produced. The comparison results are shown in Table II. As compared to Cloud Function, Fog Function can reduce more than 95% cross-node traffic and about 80% service latency by leveraging data locality. As compared to Edge Function, Fog Function can reduce 30% service latency for dealing with big entities thanks to the task migration mechanism, but it slightly introduces

TABLE II: Comparison with existing function-based programming models

Approaches	cross-node traffic (MB)		avg. service latency (ms)	
	small	big	small	big
Cloud Function	86.6	987.3	262	610
Edge Function	3.5	10.8	68	150
Fog Function	3.8	11.4	59	102

5% additional cross-node traffic. For time-sensitive services, this improvement is important. Also when considering more dynamic services, the benefit of Fog Function could be even more clear because Edge Function will suffer from workload imbalance.

VI. RELATED WORK

A. Fog/Edge Computing

There are already various fog computing or edge computing frameworks and approaches. Their programming models and orchestration mechanisms are defined at different levels, from the lower layer infrastructure level to the upper layer application level. For instance, Cloudlet [15] proposes an edge computing approach to offload computation from mobile devices to the network edge using virtual machine (VM) based cloudlets. KubeEdge [16] is an open source system extending native containerized application orchestration and device management to hosts at the edge. Telcofog [17] defines a unified fog and cloud computing infrastructure for 5G networks over distributed clouds based on both VMs and containers. These frameworks are suitable for generic standalone applications that can be hosted either in the cloud or at edges, but they lack a programming model to program the logic of data-intensive applications. To overcome this issue, many dataflow-based approaches are proposed. For example, the initial version of FogFlow [13] can program IoT services over cloud and edges based on service topology. AWS Step Function [18] is able to build distributed applications using visual workflows. These frameworks are flexible to support the function composition within a single application, but they require users to manually trigger the defined application and also the cross-applications function composition is not possible. To address these limitations, the function-based programming model is introduced by many fog computing frameworks, such as Amazon Greengrass [5], Azure IoT Edge [4], and Baidu

OpenEdge [7]. They all trigger functions per edge based on a topic-based pub/sub system. Opposed to that, our design can automatically trigger functions for user-definable data granularity based on the availability of their input data.

B. Serverless Computing

Serverless computing is emerging as a new paradigm for the deployment of cloud applications. Many of the major cloud vendors, including Amazon Lambda, Google Cloud Functions, Microsoft Azure Functions, IBM Cloud Functions, have released their serverless computing platforms. In addition, there are also many open source serverless computing frameworks, such as OpenWhisk, Kubeless, Fission, and OpenFaaS. They are all designed for the cloud environment in which computation resource and storage resource are unlimited and centralized. Some studies address the cold start issue of function provisioning [19], [20]. On the other hand, as serverless computing is applied into data-intensive applications [21], [22], database [23], video analysis [24], the data management and sharing between serverless tasks turns to be a bottleneck. Pocket [25] proposes a fast storage system for ephemeral data sharing between serverless tasks, yet it is still not able to benefit from data locality since the management of data and tasks is separated. Our design can manage data and serverless tasks jointly to achieve better efficiency in more geo-distributed and heterogeneous environments.

VII. CONCLUSION AND OUTLOOK

In this paper we take a first step into applying the serverless concept into fog computing for data-intensive IoT services. The goal is to keep the simplicity and flexibility of the Function-as-a-Service programming model for fog computing via an extended function programming model called Fog Function, meanwhile improving the efficiency of existing frameworks with two approaches: content-based discovery and context-driven orchestration. The current approach is scalable with hundreds of fog nodes, but it is necessary to decentralize the discovery and orchestration for a much larger scale. Also, in the future the context-driven orchestration mechanism can be improved to provide predictable service level objectives with some machine learning based approaches as proposed in [26].

REFERENCES

- [1] I. Stoica, D. Song, R. A. Popa, D. Patterson, M. W. Mahoney, R. Katz, A. D. Joseph, M. Jordan, J. M. Hellerstein, J. E. Gonzalez *et al.*, "A berkeley view of systems challenges for ai," *arXiv preprint arXiv:1712.05855*, 2017.
- [2] G. Ananthanarayanan, P. Bahl, P. Bodík, K. Chintalapudi, M. Philipose, L. Ravindranath, and S. Sinha, "Real-time video analytics: The killer app for edge computing," *computer*, vol. 50, no. 10, pp. 58–67, 2017.
- [3] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, 2017.

- [4] "Azure iot edge," <https://docs.microsoft.com/en-us/azure/iot-edge/>, 2019.
- [5] "Aws iot greengrass," <https://aws.amazon.com/greengrass/>, 2019.
- [6] "Edgex," <https://www.edgexfoundry.org/>, 2019.
- [7] "Openedge," <https://github.com/baidu/openedge>, 2019.
- [8] E. Jonas, J. S. Smith, V. Sreekanti, C.-C. Tsai, A. Khandelwal, Q. Pu, V. Shankar, J. Carreira, K. Krauth, N. Yadwadkar, J. E. Gonzalez, R. A. Popa, I. Stoica, and D. A. Patterson, "Cloud programming simplified: A berkeley view on serverless computing," 2019.
- [9] J. M. Hellerstein, J. Faleiro, J. E. Gonzalez, J. Schleier-Smith, V. Sreekanti, A. Tumanov, and C. Wu, "Serverless Computing: One Step Forward, Two Steps Back," in *The biennial Conference on Innovative Data Systems Research (CIDR)*, 2019.
- [10] L. Wang, M. Li, Y. Zhang, T. Ristenpart, and M. Swift, "Peeking behind the curtains of serverless platforms," in *2018 USENIX Annual Technical Conference (USENIX ATC 18)*. Boston, MA: USENIX Association, 2018, pp. 133–146.
- [11] I. Baldini, P. Castro, P. Cheng, S. Fink, V. Ishakian, N. Mitchell, V. Muthusamy, R. Rabbah, and P. Suter, "Cloud-native, event-based programming for mobile applications," in *Proceedings of the International Conference on Mobile Software Engineering and Systems*. ACM, 2016, pp. 287–288.
- [12] M. Bauer, E. Kovacs, A. Schülke, N. Ito, C. Criminisi, L. W. Goix, and M. Valla, "The context api in the oma next generation service interface," in *2010 14th International Conference on Intelligence in Next Generation Networks*, Oct 2010, pp. 1–5.
- [13] B. Cheng, G. Solmaz, F. Cirillo, E. Kovacs, K. Terasawa, and A. Kitazawa, "Fogflow: Easy programming of iot services over cloud and edges for smart cities," *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 696–707, 2018.
- [14] F. Manco, C. Lupu, F. Schmidt, J. Mendes, S. Kuenzer, S. Sati, K. Yasukata, C. Raiciu, and F. Huici, "My vm is lighter (and safer) than your container," in *Proceedings of the 26th Symposium on Operating Systems Principles*. ACM, 2017, pp. 218–233.
- [15] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," *IEEE pervasive Computing*, no. 4, pp. 14–23, 2009.
- [16] "Kubeedge," <https://github.com/kubeedge/kubeedge>, 2019.
- [17] R. Vilalta, V. López, A. Giorgetti, S. Peng, V. Orsini, L. Velasco, R. Serral-Gracia, D. Morris, S. De Fina, F. Cugini *et al.*, "Telcog: A unified flexible fog and cloud computing architecture for 5g networks," *IEEE Communications Magazine*, vol. 55, no. 8, pp. 36–43, 2017.
- [18] "Aws step functions," <https://aws.amazon.com/step-functions/>, 2019.
- [19] I. E. Akkus, R. Chen, I. Rimac, M. Stein, K. Satzke, A. Beck, P. Aditya, and V. Hilt, "SAND: Towards high-performance serverless computing," in *2018 USENIX Annual Technical Conference (USENIX ATC 18)*. Boston, MA: USENIX Association, 2018, pp. 923–935.
- [20] E. Oakes, L. Yang, D. Zhou, K. Houck, T. Harter, A. Arpaci-Dusseau, and R. Arpaci-Dusseau, "SOCK: Rapid task provisioning with serverless-optimized containers," in *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, Boston, MA, 2018, pp. 57–70.
- [21] E. Jonas, Q. Pu, S. Venkataraman, I. Stoica, and B. Recht, "Occupy the cloud: Distributed computing for the 99%," in *Proceedings of the 2017 Symposium on Cloud Computing*. ACM, 2017, pp. 445–451.
- [22] V. Ishakian, V. Muthusamy, and A. Slominski, "Serving deep learning models in a serverless platform," in *2018 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 2018, pp. 257–262.
- [23] J. Schleier Smith, "Serverless foundations for elastic database systems."
- [24] L. Ao, L. Izhikevich, G. M. Voelker, and G. Porter, "Sprocket: A serverless video processing framework," in *Proceedings of the ACM Symposium on Cloud Computing*. ACM, 2018, pp. 263–274.
- [25] A. Klimovic, Y. Wang, P. Stuedi, A. Trivedi, J. Pfefferle, and C. Kozyrakis, "Pocket: elastic ephemeral storage for serverless analytics," in *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI18)*, 2018, pp. 427–444.
- [26] M. F. Argerich, B. Cheng, and J. Fürst, "Reinforcement learning based orchestration for elastic services," *arXiv preprint arXiv:1904.12676*, 2019.



The research leading to these results has received funding from the European Community's Horizon 2020 research and innovation programme under grant agreement n° 779747.