

# FaaS\$: A Transparent Auto-Scaling Cache for Serverless Applications

Francisco Romero<sup>1</sup>  
Stanford University

Gohar Irfan Chaudhry  
Microsoft Research

Íñigo Goiri  
Microsoft Research

Pragna Gopa  
Microsoft Azure

Paul Batum  
Microsoft Azure

Neeraja J. Yadwadkar  
Stanford University

Rodrigo Fonseca  
Microsoft Research

Christos Kozyrakis  
Stanford University

Ricardo Bianchini  
Microsoft Research

## ABSTRACT

Function-as-a-Service (FaaS) has become an increasingly popular way for users to deploy their applications without the burden of managing the underlying infrastructure. However, existing FaaS platforms rely on remote storage to maintain state, limiting the set of applications that can be run efficiently. Recent caching work for FaaS platforms has tried to address this problem, but has fallen short: it disregards the widely different characteristics of FaaS applications, does not scale the cache based on data access patterns, or requires changes to applications. To address these limitations, we present FaaS\$, a transparent auto-scaling distributed cache for serverless applications. Each application gets its own cache. After a function executes and the application becomes inactive, the cache is unloaded from memory with the application. Upon reloading for the next invocation, FaaS\$ pre-warms the cache with objects likely to be accessed. In addition to traditional compute-based scaling, FaaS\$ scales based on working set and object sizes to manage cache space and I/O bandwidth. We motivate our design with a comprehensive study of data access patterns on Azure Functions. We implement FaaS\$ for Azure Functions, and show that FaaS\$ can improve performance by up to 92% (57% on average) for challenging applications, and reduce cost for most users compared to state-of-the-art caching systems, *i.e.* the cost of having to stand up additional serverful resources.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*SoCC '21, November 1–4, 2021, Seattle, WA, USA*

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8638-8/21/11...\$15.00

<https://doi.org/10.1145/3472883.3486974>

## CCS CONCEPTS

• **Computer systems organization** → **Cloud computing**.

## KEYWORDS

Cloud computing, serverless computing, function as a service (FaaS), caching

### ACM Reference Format:

Francisco Romero<sup>1</sup>, Gohar Irfan Chaudhry, Íñigo Goiri, Pragna Gopa, Paul Batum, Neeraja J. Yadwadkar, Rodrigo Fonseca, Christos Kozyrakis, and Ricardo Bianchini. 2021. FaaS\$: A Transparent Auto-Scaling Cache for Serverless Applications. In *ACM Symposium on Cloud Computing (SoCC '21)*, November 1–4, 2021, Seattle, WA, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3472883.3486974>

## 1 INTRODUCTION

**Motivation.** Function-as-a-Service (FaaS) is an increasingly popular way of deploying applications to the cloud. With FaaS, users deploy their code as stateless functions and need not worry about configuring or managing resources explicitly. FaaS shifts these responsibilities to the provider (*e.g.*, AWS Lambda, Azure Functions, Google Cloud Functions), which charges users per resource usage during their function invocations. FaaS providers build their platforms by renting and managing resources (VMs or bare-metal containers) in public clouds (*e.g.*, AWS, Azure, Google Cloud Platform). To control their costs, these providers proactively unload a function from memory if it has not been invoked for a while (*e.g.*, after 7 minutes of inactivity in AWS Lambda [50]).

Due to FaaS's stateless nature, a function invocation is not guaranteed to have access to state created by previous invocations. Thus, any state that might be needed later must be persisted in remote storage. This also applies to applications with multiple stages (often expressed as a pipeline or a directed acyclic graph of functions), with intermediate results passed between invocations. Since existing FaaS platforms

<sup>1</sup>Romero ([faromero@stanford.edu](mailto:faromero@stanford.edu)) was an intern at Microsoft Research.

typically do not allow functions to communicate directly, functions must also write these results to remote storage.

The remote storage can be object-based (e.g., Amazon S3 [2], Azure Blob Storage [37]), queues [8], or in-memory storage clusters (e.g., Redis [47], InfiniCache [56], Pocket [34]). Regardless of type, remote storage incurs higher latency and lower bandwidth than accessing local memory [26, 34]. When users have to provision in-memory storage clusters, it introduces management overhead and costs.

Given these limitations, local in-memory caching emerges as a natural solution for both speeding up access to remote data and enabling faster data sharing between functions. Prior works have considered both local and remote in-memory caching for FaaS [34, 45, 53, 56] but, we argue, have come up short in multiple ways.

First, prior work has shown that FaaS applications vary widely in their invocation frequency, e.g. 45% of applications are invoked less frequently than once per hour on average [48]. Thus, caching data for rarely-invoked applications at all times is wasteful. However, not caching data for these applications will likely produce poor performance. Prior systems implement a single cache for multiple or all applications, which given these characteristics, requires complex communication and synchronization primitives for the data of thousands of applications.

Second, in prior approaches, the cache is either fixed in size (e.g., [45, 56]) or scales only according to the computational load (e.g., [53]). These approaches work well when data access patterns are stable and working sets are smaller than the available cache space. When this is not the case, scaling the cache based on data access patterns would be beneficial. Moreover, prior works have not considered scaling as a way to mitigate the impact of accessing large data objects; these objects can take long to access as remote storage I/O bandwidth is often limited by the underlying VM/container or contention across co-located applications [26]. Emerging FaaS applications, such as ML inference pipelines and notebooks, would benefit from scaling out for increased cache space and increased I/O bandwidth to remote storage.

Third, prior caches are not entirely transparent to users, either because users need to provision them explicitly (e.g., [34, 45]) or because they provide a separate API for users to access the cache (e.g., [51, 53]). FaaS users do not want to think of data locality or manage caches (or any other resources) explicitly. In fact, a key reason for the popularity of FaaS is exactly that users are relieved from such tasks.

**Our work.** In this paper, we propose FaaS\$, an in-memory caching layer for FaaS applications that is itself serverless, *i.e.* tied to applications, auto-scaling, and transparent.

Each application is loaded into memory with its own local cache. This cache manages the data accessed by the application transparently as it runs. When the application is

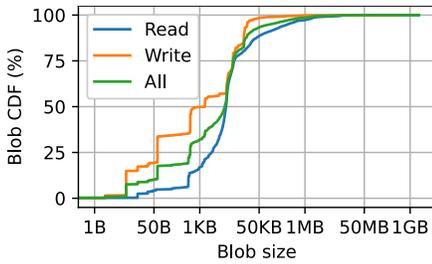
unloaded from memory, its FaaS\$ is also unloaded. This approach obviates the need for a remote in-memory cache and may reduce the overall traffic to remote storage, both of which reduce costs for users. It also means that the cache space required by rarely-invoked applications is proactively removed from memory when not needed, just as the application itself, which reduces costs for the FaaS provider. Moreover, it enables different cache replacement and persistence policies per application, and pre-fetching of the most popular data when (re-)loading each application. The latter feature can be very effective when combined with automatic pre-warming of applications, which can be done accurately and completely off the critical invocation path [48].

As in other systems, an application is auto-scaled based on the number of invocations it is receiving. Scaling out loads a new application “instance” (*i.e.*, a copy of its code) into memory, whereas scaling in unloads an instance. We refer to this as compute-based scaling. However, to match each application’s data access and reuse patterns, FaaS\$ automatically scales out the number of instances to (a) increase the fetch bandwidth from remote storage for large objects, and (b) increase the overall cache size for frequently-accessed remote data. Scale-in occurs as the need for space and/or bandwidth subsides. With multiple active instances, FaaS\$ forms a cooperative distributed cache, *i.e.* a data access can hit locally, miss locally but hit a remote cache, or miss both locally and remotely. By default, FaaS\$ offers strong data consistency, but each application can *optionally* select its own consistency and policies for scaling and eviction.

A key aspect of FaaS is that users are only charged for resources they use. As we tie FaaS\$ to applications, we expect that FaaS providers would charge only for cache accesses and space consumed by objects that were actually accessed. **Implementation and results.** We motivate FaaS\$ with the first comprehensive study of real FaaS applications from the perspective of data access and reuse. We use data collected for 2 weeks from the production workload of Azure Functions. We show, for example, that many infrequently invoked applications exhibit good temporal locality (the same data is accessed across relatively rare invocations), whereas spatial locality in large objects is high (if any byte from such an object is accessed, the rest of it should be pre-fetched).

We implement FaaS\$ for Azure Functions. To show it enables new applications that are not efficiently supported by current FaaS platforms, we implement an ML inference pipeline and a Jupyter notebooks [32] server stack that runs unmodified notebooks in a serverless environment.

Our experiments with these applications evaluate FaaS\$’s caching and scaling policies. Our results show that FaaS\$ can improve their performance by up to 92% (57% on average), and reduce costs for most users compared to state-of-the-art



**Figure 1: Distribution of the size of accessed blobs. 80% of accessed blobs are smaller than 12KB.**

FaaS caching systems, *i.e.*, the cost incurred by the provisioning of additional serverful resources.

**Contributions.** In summary, our main contributions are:

- We characterize the data access patterns of Azure Functions. We release a sanitized version of this dataset at [6].
- We design and implement FaaS\$, a transparent auto-scaling cache for FaaS applications that is itself serverless.
- We propose scaling policies for FaaS\$ to increase instance bandwidth and overall cache size based on data access patterns and object sizes.
- We show that FaaS\$ broadens the scope of applications that can run on FaaS with near-native performance, including ML inference pipelines and Jupyter notebooks.

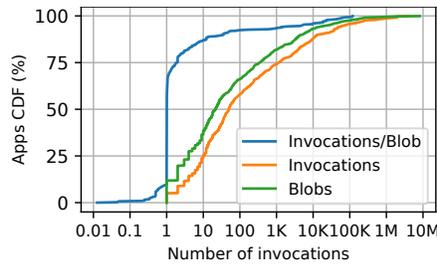
## 2 FAAS APPLICATIONS AND CACHING

This section characterizes the invocation and data access patterns of applications running on a production FaaS offering. An *application* is a collection of logically-related functions. Each instance of an application gets a set of resources (*e.g.*, memory in a VM or container) that is shared by its functions. We focus on data accesses, as several prior works have focused on code accesses to optimize cold-start latencies [1, 13, 18, 40] or reducing the number of cold-starts [48]. We also discuss the limitations of current FaaS platforms for existing and emerging applications.

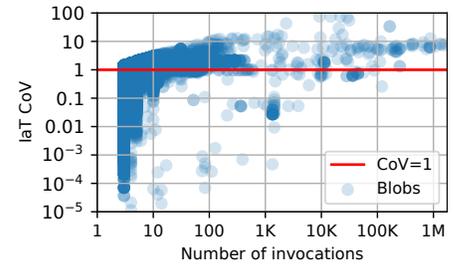
### 2.1 Characterizing Current Applications

We deployed a new version of the Azure Functions runtime to characterize blob storage accesses. We randomly sampled a small percentage of applications that access remote blob storage over HTTPS over a 14 day period (November 23<sup>rd</sup> to December 7<sup>th</sup> 2020). This log spans 28 geographical regions, includes 855 applications from 509 users, and 33.1 million invocations with 44.3 million data accesses. 77.3% of accesses are reads and the rest are writes. The applications are written in multiple languages, including C#, Node.js, and Python. With this paper, we are releasing an anonymized version of this log at <https://github.com/Azure/AzurePublicDataset>.

**Data size.** The log includes accesses to 20.3 million different objects with a total size of 1.9 TB. Figure 1 shows the distribution of the size of blobs accessed. 80% of blobs are



**Figure 2: Application CDF of the number of function invocations per unique blob.**



**Figure 3: IaT CoV and number of invocations for each blob accessed more than three times. CoV of 1 is Poisson arrival.**

smaller than 12KB and more than 25% are smaller than 600 bytes. However, there are also many large blobs; a few as large as 1.8GB. The objects read tend to be larger than the ones written. While the aggregate bandwidth to backend storage is usually high [49], the prevalence of small objects exacerbates the impact of storage latency.

**Data accesses and reuse.** Figure 2 shows the CDF of the ratio of the number of invocations an application made to the number of unique blobs accessed. Most applications access a single, different blob per invocation ( $\text{invocation}/\text{blob}=1$ ). Roughly 11.0% of applications access more than one blob per invocation ( $\text{invocation}/\text{blob}<1$ ). More than 32.0% of the applications access the same blob in more than one invocation and 7.7% in more than 100 invocations. One application accesses the same blob in more than 10,000 invocations.

Around 11.8% of the applications access the same blob across all invocations, 66.1% access less than 100 different blobs, 93.0% access less than 10,000 different blobs, and one accesses more than 8 million different blobs. Even though there are 44.3 million accesses, only 20.3 million are first accesses. Overall, the applications accessed 2.6 TB of data while the corpus of unique data is 1.9 TB. *If we were able to cache the already accessed data, we could save up to 27.0% of traffic and 54.3% of the accesses to remote storage.*

Data sharing across applications and users (not shown) is extremely uncommon. 99.7% of the blobs are not shared across applications, only 0.02% of blobs are shared across regions, and only 16 blobs across users.

**Temporal access pattern.** Figure 3 shows the temporal access patterns for each blob accessed by an application. The X-axis is the number of function invocations that read/wrote the blob and the Y-axis shows the coefficient of variation (CoV) for the inter-arrival time (IaT) of those invocations. Each point represents a blob with more than 3 accesses (there is no IaT CoV otherwise). A CoV of 1 suggests Poisson arrivals, values close to 0 indicate a periodic arrival, and values larger than 1 indicate greater burstiness than Poisson. Clearly, accesses to a large percentage of blobs are bursty.

**Access performance.** We observe writes are usually faster than reads, since writes are buffered and do not require

Data size	Large	Large	Large	Large	Small	Small	Small	Small
App. invoc. freq.	Frequent	Rare	Frequent	Rare	Frequent	Rare	Frequent	Rare
Data reuse	Low-med	Med-high	Med-high	Med-high	Low	Med	Low-med	Low
Example application	Video streaming	Jupyter notebook	ML inference (pipeline)	Distributed compilation	Log aggregator	Software unit tests	ML inference (single model)	IoT sensors

**Table 1: FaaS application spectrum. Applications vary in data size, invocation frequency, and data reuse.**

Features/Properties	Pocket [34]	InfiniCache [56]	Locus [45]	Faasm [51]	Cloudburst [53]	OFC [39]	FaaS†
Cache location	Storage cluster	Separate function	Redis cluster	Host cache	VM cache	RAMCloud server	Invoked function
Cached data management	Independent of app	Unloaded w/ app	Unloaded w/ app				
App transparency	Get/Put	Get/Put	No changes*	Custom API	Custom API	No changes	No changes
User configuration	Hints	Num instances	Redis	None	None	None	None
Data consistency	None	None	None	Supported	Supported	Supported	Supported
Object pre-warming	None	None	None	None	None	None	Automatic
Compute cache scaling	N/A	N/A	N/A	Dynamic	Dynamic	Dynamic	Dynamic
Cache size scaling	N/A	None	None	None	None	Dynamic	Dynamic
Cache bandwidth scaling	N/A	None	None	None	None	None	Dynamic

**Table 2: Existing caching systems and their properties. Green indicates supported/enabled. Tan indicates limited support, or limitations in what can be enabled. Red indicates not supported/enabled. For Locus, \* indicates the application transparency is limited to applications with a shuffle operation.**

persisting data across all replicas synchronously. Reads are slower as we have to wait for the storage layer to deliver all data. Furthermore, smaller blobs have a lower throughput as they cannot amortize the overhead of the initial handshake. **Diverse invocation patterns.** Consistent with prior work [48], our log shows that 80% of functions have less than one invocation per minute on average and less than 15% of applications account for 99% of the invocations. This heterogeneous invocation behavior is a challenge, as caching data for rarely-invoked applications can be wasteful but is necessary, as it affects a large number of users.

**Takeaways and requirements.** Our characterization indicates that many FaaS applications exhibit data reuse: more than 30% of them access the same data across invocations. This suggests that caching can be effective for them. Moreover, the characterization shows a wide spectrum of accessed data sizes and invocation frequencies. Accessed data sizes span almost 9 orders of magnitude (from several bytes to GBs), *i.e.* large objects cannot be overlooked. Function invocation frequencies also span almost 9 orders of magnitude, *i.e.* rarely-invoked applications cannot be overlooked.

Table 1 illustrates the spectrum of data sizes, invocation frequencies, and reuse along with some example applications. For instance, distributed compilation of the Chromium browser requires accesses to hundreds of MBs, but happens only a few times per day using a framework like gg [19]. Data reuse across compilations is high since the codebase does not change fast. In contrast, serving an IoT sensor involves a small dataset, rare invocations, and low reuse.

We draw three key requirements for a serverless caching layer. First, it should ensure that *data with good temporal*

*locality is cached* and reused across invocations (R1). Second, the caching layer should target *both frequently- and rarely-invoked applications* (R2). It should optimize for data reuse for frequently-invoked applications, while it should have the ability to pre-warm frequently-accessed objects for rarely-invoked applications. Finally, the caching layer should *accommodate large objects and exploit spatial locality* for them (R3). These requirements must be satisfied for applications written across various programming languages.

## 2.2 Existing Caching Systems Limitations

Table 2 lists the characteristics of several caching systems. **Caching is managed independently of each application.** Except for OFC, all the systems listed include a separate caching or storage infrastructure that is shared by multiple applications. (Cloudburst and Faasm also have a shared cache on the same hosts/VMs that run the functions.) Because of this, either users must manage the extra servers or cache state is left behind when applications are unloaded from memory. In the former case, user costs and management overheads are greater, whereas in the latter the FaaS provider costs are higher. Thus, a serverless caching layer *should be tied to each application, so that its code and data can be loaded/unloaded based on the application’s invocation pattern* (R4).

**Need system configuration or application changes.** Pocket, InfiniCache, and Locus rely on user configuration to maximize performance and minimize cost. For example, InfiniCache users must statically set the number of data shards and functions to store them, whereas Locus users must configure a Redis cluster. Faasm and Cloudburst rely on custom APIs to give applications control over the consistency of their data, and to pass messages between functions. To retain the

simplicity of the FaaS paradigm, the caching layer *should be transparent and not require changes to the application* (R5).

Cloudburst additionally requires its key-value store (Anna [57]) to track all objects residing in all caches. Managing this much metadata can lead to scaling limitations and extra costs. The caching layer can mitigate these concerns by *minimizing metadata management for each cache instance* (R6).

**Compute-based scaling only.** Finally, while some existing systems do not dynamically scale their caches (Pocket, InfiniCache, Locus), others do so based solely on the amount of offered computational load, *i.e.* number of function invocations (Faasm, Cloudburst). OFC scales based on the computational load and predicted memory usage of cached objects. However, it limits its caching to objects smaller than 10MB. As applications become more complex and data-heavy, data access characteristics like large working sets or large objects will gain in importance. Thus, the caching layer *should scale compute (as the application’s offered load varies), cache size (based on the data reuse pattern), and bandwidth to remote storage (based on the object sizes being accessed)* (R7).

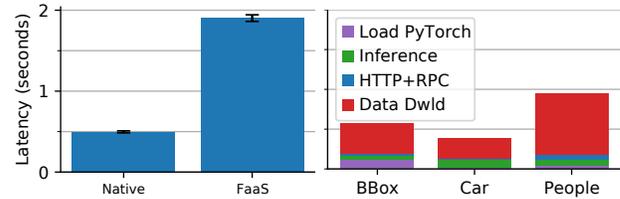
### 2.3 Enabling New FaaS Applications

Current FaaS platforms limit the set of applications that can be efficiently run. Next, we describe two challenging ones.

**ML inference pipeline.** Many applications across several domains (*e.g.*, health care [29] and advertisement [23]) depend on ML inference for various prediction tasks. ML inference load patterns can vary unpredictably [46, 58], which makes FaaS on-demand compute and scaling an ideal match for serving inference queries. However, ML inference applications require low-latency prediction serving (<1 s) [22, 46]. For example, AMBER Alert [36] responders may use an application to perform car and facial recognition. This can be deployed on a FaaS platform as a pipeline of ML models. For each input image, a request is first sent to a bounding box model function to identify and label all present objects. The labeled image is uploaded to a common data store to trigger the car and people recognition models. Both recognition functions upload their outputs to the common data store.

To assess whether FaaS can meet low-latency requirements, we ran the AMBER Alert pipeline natively on a local VM, and on Azure Functions with remote storage. Figure 4a shows it is up to 3.8× slower running on FaaS versus natively, while Figure 4b shows that the main reason is that 73.1% of the time is spent loading the models. The inefficiency of the storage layer makes it impossible for the FaaS platform to run this application with sub-second latency.

**Jupyter notebooks.** Jupyter notebooks [32] are often used for data science tasks such as data processing, ML modeling, and visualizing results [43]. They are typically run by defining code in cells and interactively executing each cell. Jupyter notebooks are typically backed by statically-allocated VMs.



(a) Native VM versus prod FaaS. (b) FaaS execution breakdown.

**Figure 4: (a) Latency of AMBER Alert pipeline on a native VM versus Azure Functions. Native VM does not include the PyTorch (~700MB) loading time. It is up to 3.8× slower to run on FaaS. (b) Model run times (BBox is a bounding box model). Data movement from/to remote storage dominates.**

Depending on how often a notebook is run, the VMs may sit idle for long periods. This is expensive for users and wasteful for service providers. Furthermore, the amount of parallelism and compute needed for *each cell’s task* can vary. Akin to ML inference, this variability makes FaaS a strong fit.

To test its performance, we ported Jupyter to run on Azure Functions — an application we term *JupyterLess*. Each cell is executed as a function invocation and the state between functions is shared through an intermediate storage layer. We compare the execution time of summing a single 350MB DataFrame column partitioned into 10 chunks with JupyterLess to running on a native VM. JupyterLess is 63× slower than native VM execution as downloading the intermediate state and DataFrame column from remote storage dominates the execution time. Thus, JupyterLess cannot be run interactively on existing FaaS frameworks.

### 3 FaaS\$ DESIGN

We design FaaS\$ as a transparent auto-scaling cache that meets the requirements from Section 2. FaaS\$ caches objects accessed during a function execution so they can be reused across invocations (R1). It is built into the FaaS runtime with no external servers or storage layers, so it can be transparently tied to an application (R4, R5) written in any of the various supported languages. This also prevents data or resources from being persisted beyond an application’s lifetime, which can incur higher costs for FaaS providers. When an application is unloaded from memory, FaaS\$ collects metadata about the cache objects, and uses it to pre-warm the cache with frequently accessed objects when the application is re-loaded into memory. This is especially important for applications that are rarely-invoked (R2).

FaaS\$ scales along three dimensions (R7): (a) based on an application’s invocations per second (*compute scaling*), which benefits applications that are frequently-invoked; (b) based on the data reuse pattern (*cache size scaling*), which is beneficial for applications with large working sets whose objects are continuously evicted; (c) based on the object size (*bandwidth scaling*), which is beneficial for applications

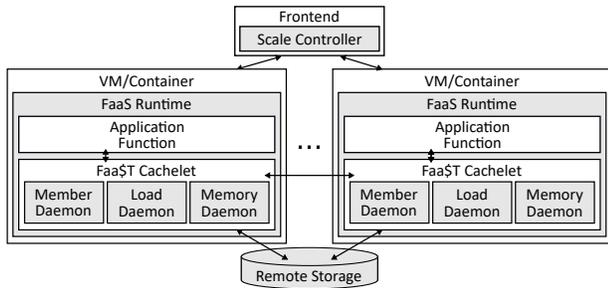


Figure 5: FaaS’s architecture diagram.

that access large objects ( $\geq 10\text{MB}$ ) and are limited by the I/O bandwidth between the application instance and remote storage (R3). While an application is loaded, FaaS efficiently locates objects across instances using consistent hashing without the need for large location metadata (R6).

### 3.1 System architecture

Figure 5 shows the architecture of a FaaS platform with FaaS. Each application instance runs in a VM or container that contains the FaaS runtime and the code for the application functions. FaaS instances, which we refer to as *cachelets*, are a part of the runtime, caching data in memory. Each application instance has one corresponding FaaS cachelet. In addition, FaaS forms a cooperative distributed cache; an application’s FaaS cachelets communicate directly to access data when necessary (Section 3.2). Similar to Faasm, FaaS maintains a single copy of cached objects, which improves memory efficiency compared to existing systems [39, 53, 56].

We designed FaaS to be per-application as opposed to a shared cache. While shared caches can multiplex applications to improve resource utilization, they suffer from the following drawbacks. First, a shared cache requires complex communication and synchronization primitives for the data of thousands of applications (compared to tens of instances for a single application with its own cache). This makes it difficult to implement per-application management policies and provide transparency without custom APIs [51, 53], especially given the diversity of application characteristics and requirements (Section 2.1). Second, a shared cache with traditional eviction policies (e.g. LRU) can lead to severe unfairness among applications [44].

To find the location of an object, a FaaS cachelet interacts with the *Membership Daemon*, which determines the object’s “owner” based on the current number of cachelets. The owner is responsible for downloading/uploading the object from/to remote storage. The *Load Daemon* collects cached object metadata, and uses it to decide what data objects to pre-warm when an application is loaded (Section 3.3). To prevent interference with an application’s heap memory usage,

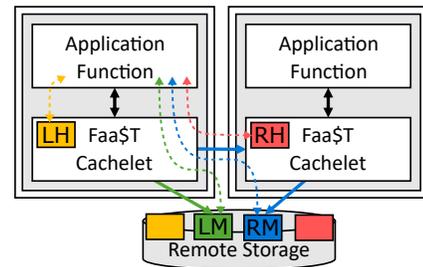


Figure 6: Reads in FaaS: local hit (LH), remote hit (RH), local miss (LM), remote miss (RM). Solid lines indicate communication between the application, FaaS instances, and remote storage. Dashed lines indicate data movement.

the *Memory Daemon* monitors both function and cachelet memory consumption. Finally, the *Frontend* load-balances requests across the running application instances, and the *Scale Controller* adds and removes instances based on metrics provided by the FaaS runtime (Section 4).

### 3.2 Accessing and caching data

**Reads.** Figure 6 shows the four ways to read data. A *local hit* finds the data cached in the local FaaS cachelet. A *local miss* occurs when the local FaaS cachelet is the owner for the object and does not currently cache the data. The cachelet directly fetches the data from remote storage. A *remote hit* occurs when the data misses the local cachelet but is found in the owner’s cache. Finally, a *remote miss* occurs when the access misses both the local cache and the owner’s cache. The owner fetches the data from remote storage and caches it locally. In all cases, FaaS cachelets cache objects locally, even if they are not the owners, for performance and locality. Thus, a popular object will incur at most one remote hit per cachelet, and local hits thereafter (besides the optional consistency version check, described below).

FaaS uses consistent hashing to determine object ownership. We choose consistent hashing because (a) it avoids having to track object metadata (e.g. list of objects in each instance), and (b) it reduces rebalancing as instances are added/removed: on average, only  $numObjects/numInstances$  need to be remapped [5, 33]. To maintain transparency, the object namespace is the same as that used by the remote storage service. This design choice also enables efficient communication between cachelets, which as observed by prior work [14, 19, 31] is beneficial for applications such as ML training.

**Writes.** When the application needs to output data, FaaS writes through to the owner cache. The instance executing the function sends the data to the owner cache, which then writes it to remote storage. This guarantees that the owner always has the latest version that the application has written. By default, the write happens synchronously to the owner and synchronously to remote storage. This offers high fault

Write target	Write mode	Read target	Performance	Consistency	FT
Storage	Sync	Storage	Low	Strong	High
Owner	Sync	Owner	Medium	Strong	Medium
Owner	Async	Owner	Medium	Eventual	Medium
Owner	Sync	Local	High	Weak	Medium
Local	Sync	Local	High	None	Low
Local	Async	Local	High	None	Low

**Table 3: Performance, consistency, and fault tolerance (FT) for different settings. By default, FaaST writes to storage synchronously, and reads the version from storage (first row).**

tolerance while trading off performance, since applications must wait until the write completes before proceeding with their execution. Applications can *optionally* configure FaaST to write asynchronously or not write to remote storage at all. Because FaaST is tied to each application, different applications can use different policies at the same time.

**Consistency.** Table 3 summarizes FaaST’s read/write settings, and the performance, consistency, and fault tolerance (FT) they achieve. By default, when reading an object, FaaST first verifies the version in the cache matches the one in remote storage. From our measurements with Azure Blob Storage, verifying the version from remote storage takes less than 5ms (compared to 10ms and 500ms for retrieving a 4B object and 35MB object, respectively). If the version matches (the common case), no object is retrieved. This verification, combined with synchronous writes to remote storage, offers strong consistency (first row of Table 3). This default provides the same fault tolerance with better performance than the production FaaS platform.

Given the latency of verifying the version, some applications may be willing to trade off consistency for performance (e.g., ML inference). For those applications, FaaST can read any cached version and write asynchronously to remote storage. This weakens its fault tolerance, and provides only eventual consistency. Applications can also completely skip writing to remote storage and rely on the distributed cache. In Section 6.6, we quantify the performance and consistency impact of these settings.

### 3.3 Pre-warming data into FaaST

To *pre-warm* future cachelets, FaaST records metadata about the cache, off the critical path, when the Frontend unloads the application. This includes the size of each cached object, its version, its number of accesses of each type (e.g., local hit or if it was produced as an output), and its average inter-arrival access time. We timestamp each metadata collected with the unload timestamp to capture the state history of the cachelet. As we describe next, this is necessary for applications that are rarely-invoked (e.g., once per hour), since their data access pattern cannot be determined by a single invocation.

FaaST needs to decide when to load an application into memory. For this, the Frontend leverages a previously-proposed hybrid histogram policy [48]. The policy tracks the idle times between invocations of an application in a histogram. When the application is unloaded, the Frontend uses the histogram to predict when the next invocation is likely to arrive, and schedules the reload of the application for just (e.g., 1 minute) before that time. Our approach would work with any other cold-start prevention policy as well.

At this point, FaaST needs to decide what data objects should be loaded into the new cachelet. To do so, it collects and merges the metadata across all cachelets over a pre-set period of time. The period of time is based on the application’s invocation frequency, which is determined using the hybrid histogram policy. Next, FaaST determines the objects to be loaded using the following two conditions. First, if an object’s local or remote cache hit rate is greater than a threshold, the object should be loaded. This indicates that the object has temporal locality. Second, if an object is accessed more than once across the merged metadata, the object should be loaded. This benefits rarely-invoked applications by loading objects accessed *across* unload/load periods (e.g., an ML inference application’s model and labels). Once the objects to be loaded are determined, the FaaST cachelet loads the objects that it owns based on consistent hashing (Section 3.2).

To avoid competing with on-demand accesses, FaaST pre-warms the cache only when the application is not executing, *i.e.* before an invocation arrives or right after a function execution ends. If we cannot avoid a cold-start invocation, the only data that is loaded into the cache is its inputs.

For the AMBER Alert pipeline described in Section 2.3, pre-warming improves performance by 58% over the hybrid histogram policy and 74% over a cold-start. We evaluate its effectiveness for applications run in production in Section 6.2.

### 3.4 Evicting application data from FaaST

The memory capacity of each application (and thus FaaST) cachelet is set by the provider (typically a few GBs). Cachelets do not consume any memory beyond that allocated to the application. Each Memory Daemon monitors the memory usage of the function and the cachelet. When the memory consumed by the function (*i.e.*, heap memory) and the cachelet (*i.e.*, cached objects) is within a small percentage (*i.e.*, 10%) of the application’s total memory capacity, it evicts objects.

Eviction policies are often designed to cover the broad set of applications that can run on the platform [9–11]. Instead, as FaaST is tied to an application, it can use *per-application eviction policies*. Hence, the eviction policy can be kept simple and tailored to an application’s data access pattern.

We implement two policies that we expect will work well for many applications. The first (default) policy is a simple

Least-Recently-Used (LRU) policy that prioritizes the eviction of objects that are not owned by the evicting cachelet. Only after these objects are evicted, does the policy consider owned objects in LRU order. The second policy targets objects that are larger than a threshold (e.g., 12KB) and not owned by the evicting cachelet. If there are not enough of these objects, the policy evicts non-owned objects smaller than the threshold. If we still need more capacity, we evict owned objects that are larger than the threshold, before resorting to LRU for the remaining ones. In both eviction policies, targeting non-owned objects first increases remote hits, but also minimizes local and remote misses.

Each of these eviction policies fits our emerging applications nicely. ML inference applications that exhibit high invocation rates (e.g., frequently used recommendation models [24]) can quickly fill up a FaaS cachelet’s capacity with invocation inputs (e.g., images) and outputs (e.g., labeled objects). Across invocations, only the model and labels are typically reused; inputs and outputs change each time. Thus, for ML inference and similar applications, the LRU policy is sufficient, since the inputs and outputs will be evicted when the cachelet’s capacity reaches its limit.

JupyterLess data objects can be classified into two types: (a) small objects that maintain the notebook’s state (e.g., a dictionary object) and (b) larger objects that are used for data analysis (e.g., a DataFrame). A notebook’s state is typically reused across invocations, and should thus be cached as much as possible. Larger objects are reused less frequently and can be replaced more aggressively. Thus, the second policy (size-based) is appropriate.

FaaS allows for future eviction policies beyond the ones described above. For example, objects can be given a time-to-live (TTL) and get evicted when the TTL expires.

### 3.5 Charging for FaaS

When using FaaS, we expect FaaS providers to charge users only for the memory of the accessed data and not all the cached objects. FaaS providers should also not charge for pre-warming metadata in the same manner that they do not charge for function metadata (e.g. function registration).

## 4 SCALING FAAS

FaaS platforms typically include a Scale Controller that scales applications in/out (Figure 5). As it is part of the front-end component, the Scale Controller monitors the end-to-end performance and the load offered to each application. It also periodically queries the FaaS runtime running each application instance for a vote on how many more instances to add: a positive number means a vote to scale-out and a negative number means a vote to scale-in. Based on the information for an application, it makes a scaling decision and effects it. FaaS extends this mechanism by including cache-specific

metrics when deciding on how to vote. We also extend the Scale Controller to accept unrequested votes when scaling is needed immediately. When the controller adds or removes an application instance, FaaS reassigns the objects’ ownership using the Membership Daemon’s consistent hashing.

FaaS has three types of scaling:

**Compute scaling.** FaaS platforms scale the number of application instances based on its rate of incoming requests, its number of in-flight requests (queue sizes), and/or its average response time. Degrading performance, high request rates, or long queues cause scale-out; the opposite causes scale-in. Since every application instance includes both compute and caching resources, this traditional way of scaling is sufficient.

**Cache size scaling.** FaaS also scales to match the application’s working set size. For example, a JupyterLess notebook performing data-intensive operations may not fit the entire working set in the cache, leading to a high eviction rate. To address this, each cachelet tracks the number of evictions of each locally-cached object and votes to scale out by one instance, if any object has been evicted more than once since the last controller query. If no object has been evicted more than once but there is still substantial cache access traffic, FaaS votes to do nothing (add 0 instances). It votes to scale in by one instance when the frequency of accesses is low.

Many existing caching systems statically allocate resources and either cannot auto-scale their capacity as the amount of data accessed varies or require application hints to do so. OFC uses per-application machine learning models to achieve the same dynamic cache size scaling, which requires frequent retraining and mechanisms to prevent application “out-of-memory” failure.

**Bandwidth scaling.** FaaS also supports applications with large input objects. For such objects, FaaS equally partitions the download from remote storage across multiple cachelets to (a) create a higher *cumulative* I/O bandwidth to remote storage, and (b) exploit the higher communication bandwidth *between* instances ( $BW_{Inst}$ ) compared to the bandwidth between each instance and remote storage ( $BW_{BS}$ ).

When a FaaS cachelet receives an object access, it iteratively computes the data transfer latency,  $T_{DR}$ , for a number of instances  $N$  (starting at the current number) and the object size  $S$  using the following formula:

$$T_{DR} = T_{Load} + S/N \times 1/BW_{BS} + (S - S/N) \times 1/BW_{Inst}$$

where  $T_{Load}$  is the instance loading latency. FaaS periodically profiles  $T_{Load}$ ,  $BW_{BS}$ , and  $BW_{Inst}$  to account for variations in the network and the remote storage bandwidths. The iterative process stops at the  $N$  where  $T_{DR}$  changes by less than 10% or when  $T_{DR}$  increases between iterations. If the resulting  $N$  is greater than the current number of instances, the FaaS cachelet *immediately* contacts the controller to scale out to  $N$ . FaaS then waits for the new instances to be

created (by checking the Membership Daemon) and sends each of them a request to download a different byte range of size  $S/N$ . As scale-in is not as time-critical, FaaS\$ does it through periodic voting (when queried by the controller) as the number of object accesses becomes small.

We find that bandwidth-based scale-out is worthwhile for objects on the order of hundreds of MB (Section 6.5); this will become smaller as cold-start optimizations continue to appear [1, 13, 18, 40, 55].

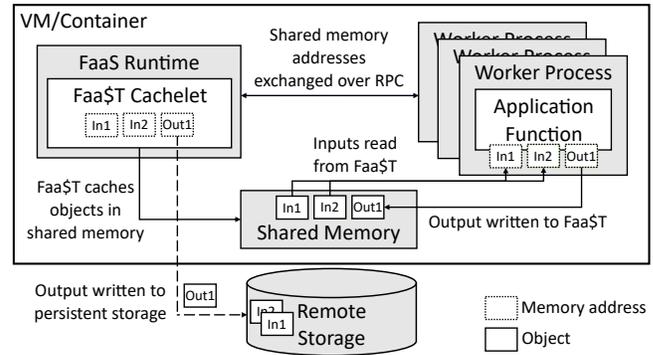
**Handling conflicting scaling requests.** The scaling policies work concurrently, so there may be scenarios where they make conflicting scaling requests to the Scale Controller. For example, compute scaling may want to scale out, while cache size scaling may want to scale in. When there are conflicting votes, the controller scales out if *any* policy determines that scale-out is needed. It scales in if *all* policies suggest scale-in will not hurt. This is similar to the approach taken by existing systems for right-sizing storage clusters [34].

**Idle function computation resources.** When instances are added for cache size or bandwidth scaling, their computation resources can be wasted. FaaS\$ minimizes resource waste by scaling in when the frequency of accesses is low. Providers can also leverage techniques for running low-priority tasks (e.g., analytics jobs) on unused resources [61].

## 5 IMPLEMENTATION

**Production FaaS platform.** We implement FaaS\$ for Azure Functions, and have open-sourced the bulk of it [15, 16]. In Azure Functions, an application comprises one or more functions. Each function defines its data bindings, which FaaS\$ uses to transparently load and manage objects: trigger (e.g., HTTP request), inputs (e.g., blob), and outputs (e.g., message queue). Users *optionally* set FaaS\$ policies (scaling, eviction, and consistency) using simple application-specific configurations during deployment. Users can also mark objects that should not be cached by FaaS\$.

As we show in Figure 7, an application instance executes in either a VM or a Docker container, and includes the FaaS runtime and function-execution workers. Upon receiving incoming requests (e.g., as a result of an incoming HTTP request, a new blob being created), the runtime collects the requested input bindings and invokes the function in a worker while passing the appropriate arguments to it. When the worker finishes executing the function, it replies to the runtime with the produced output(s) and the runtime processes them (e.g., writes a blob to remote storage or writes to a message queue). If there are multiple concurrent invocations, more worker processes can be spawned on the same instance to execute them in parallel. Azure Functions leverages Azure Blob Storage as its remote storage service.



**Figure 7: FaaS\$ integrated with the FaaS platform. Runtime and workers point to the same shared memory objects.**

As Figure 5 shows, a Frontend component handles HTTP requests and does compute scaling. We extend this component to implement bandwidth-based scaling (Section 4).

**Caching data.** We implement the core of FaaS\$ in the runtime (C# code) with minimal changes to the workers (Python and Node.js). In the original design, the runtime and workers exchanged control and data messages over a persistent RPC channel. FaaS\$ replaces the data messages with a shared memory area, while keeping control messages over RPC. The shared area is also where FaaS\$ caches data. Data communication between the FaaS\$ cachelets and the workers happens via passing shared memory addresses, reducing end-to-end latency. In addition, unlike existing systems that need to maintain duplicate object copies [39], using shared memory reduces the memory footprint by only keeping a single object copy [16]. The workers managed by the same runtime share the cached objects.

When the runtime prepares input data bindings before invoking a function, FaaS\$ intercepts them and checks the cache first (Section 3.2). When a function produces an output, FaaS\$ caches it for future invocations. This cache write triggers any functions that have the newly added object as their trigger binding. This improves latency for applications that rely on writing intermediate outputs to external sources (e.g., blob storage) to trigger subsequent functions.

We support applications written in C#, Python, and Node.js. Supporting other languages would require minimal changes. We use the shared memory APIs already available in most languages for both Linux and Windows. When we run applications in containers (vs VMs), we share (setting up permissions) the cache space across containers.

**Distributed cache.** Each runtime instance saves some metadata about the running application in a blob from remote storage. We store the FaaS\$ membership information in this blob and the cachelets periodically heartbeat their state there. The consistent hashing algorithm uses SHA256 for hashing and 100 cachelet replicas for load balancing. More replicas

did not improve load balancing and increased the ownership lookup time. Fewer replicas created ownership hot spots.

Since Azure Functions already uses HTTP for communication between its components, we use this interface to exchange data between FaaS cachelets. We evaluated other approaches like RPC (with Protocol Buffers [42]) but the improvements were negligible and the complexity of maintaining a new channel would offset them. We could also leverage RDMA but have not experimented with it.

**Other platforms.** FaaS’s design and implementation is extensible to FaaS platforms beyond Azure Functions. FaaS directly applies to the equivalent components of most platforms (e.g., runtime, worker processes, Scale Controller). All platforms use the concept of triggers and interact with external data services. However, not all of them use bindings to map the data but rely on libraries to explicitly access it inside the function body. We would need to extend these libraries (e.g., Boto3 [3] in AWS Lambda) to interact with FaaS and look-up the cache before accessing the remote storage.

## 6 EVALUATION

### 6.1 Methodology

**Comparison points.** We perform two types of comparisons. The first is an analysis of running the application logs from Section 2.1 on top of FaaS. This allows us to show the improvements these applications would get with FaaS. The second evaluates the four access scenarios that functions may encounter: objects are accessed through local hits (LH), local misses (LM), remote hits (RH), or remote misses (RM).

We compare FaaS against six baselines: (a) a large, local VM where all accesses are local and there are no function invocation overheads (*Native*); (b) Azure Functions without FaaS (*Vanilla*) that accesses all objects from remote object storage. Its performance is equal to FaaS LM; (c) InfiniCache (*IC*) [56] that we approximate by statically configuring FaaS to use only remote instances. Its best case performance is equal to FaaS RH; (d) Cloudburst’s caching layer (*CB*) [53]. Its best case performance is equal to FaaS LH; (e) *Pocket* [34], approximated with a manually managed Redis VM with all data available at memory speed (no Flash accesses); and (f) a commercial Redis service, Azure Cache for Redis (*Redis service*). Azure Cache for Redis is the offering that matches our VM size in memory and network bandwidth. It is akin to what is used by Locus [45]. Data is stored and accessed from Redis as opposed to remote object storage.

**Applications.** We use the two applications from Section 2.3: ML inference and JupyterLess notebooks. We use application latency and cost as primary metrics. For each experiment, we report the mean and standard deviation of 3 runs.

For ML inference, we evaluate both single model inference and inference pipelines. For single model inference, we use

two separate models that differ in latency and resource footprints: SqueezeNet [28] (5MB), and AlexNet [35] (239MB). For the inference pipeline, we evaluate the AMBER Alert pipeline of Section 2.3; the output of the bounding box model (MobileNet Single-Shot Detector [27], 35MB) is fed into people (ResNet50 [25], 97MB) and car recognition (SqueezeNet, 5MB) models. In all inference cases, functions access an input image, the model, and the class labels (a text file).

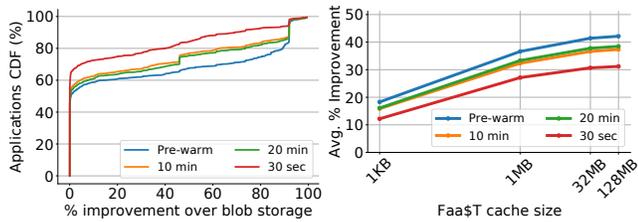
For JupyterLess, we use five notebooks: (a) single message logging (No-Op); (b) summing a 350MB DataFrame column; (c) capacity planning with data collection and plotting; (d) FaaS characterization of Section 2.1; and (e) counting up to 1K. The function data objects consist of the notebook state after each cell’s execution, which is stored in JSON format.

**Experimental setup.** Each application instance is a single VM. The default instance type (used in Sections 6.3 to 6.7) is a Standard\_DS4\_v2 [7] Azure VM with 8vCPUs, 28GiB of DRAM, and 750MB/s network bandwidth. It runs Ubuntu 18.04 with 5.4.0 kernel on Intel Xeon E5-2673 CPUs operating at 2.40GHz. I/O bandwidth to remote storage (Azure Blob Storage) is lower at 90MB/s for large objects. In our production setting, the VMs are pre-provisioned: an application instance cold-start involves loading and deploying the serverless runtime together with the application code.

**Cost model.** We derive user costs following the common pricing by FaaS providers. Function invocations are charged for the time and resources they take ( $\$/GB\text{-s}$ , order of  $10^{-5}$ ), while VMs are charged for their lifetime ( $\$/s$ , order of  $10^{-1}$ ). We assume *Native* and systems with additional resources are statically provisioned the whole time. Specifically, we charge for extra resources whenever the caching or storage system is external to the FaaS platform (i.e., Pocket, Redis, Redis Service) or specialized for FaaS in some way (Cloudburst). Except for Redis service, we charge these systems for one extra VM of the same instance size as the default application instance. The VMs are charged their on-demand prices. For Redis service, we provision per service class and charge the class’s price. The additional resource costs can be amortized by multiple applications sharing the same resources. Vanilla, InfiniCache, and FaaS use existing commodity storage (e.g., AWS S3, Azure Storage), so we do not charge them for extra resources. We also add the cost of remote storage data transfer ( $\$/op$ , order of  $10^{-6}$ ) to FaaS LM and RM, and Vanilla.

### 6.2 FaaS with production applications

We simulate the end-to-end performance of the FaaS applications from Section 2.1. Our simulator uses the default policies for consistency (synchronous writes to owner, synchronous writes to remote storage, read version from remote storage) and eviction (LRU), and implements the scaling policies described in Section 4. To model FaaS’s access latencies, we measured read and writes latencies for 1B to 2GB



**Figure 8: CDF of improvement over blob storage for a 128MB cache (left) and average improvement as the size of FaaS\$ varies (right) with the application logs from Section 2.1. Times represent different unload periods with pre-warm disabled. FaaS\$ with pre-warm improves performance by 50% or more for ~35% of applications, and has an average improvement of over 40% for a 128MB cache.**

object sizes using our FaaS\$ implementation described in Section 5. We vary the size of FaaS\$ from 1KB to 128MB; larger cache sizes showed no further improvement. We also vary the unload period and show how it affects performance when FaaS\$ cannot pre-warm frequently-accessed objects.

Figure 8 shows the CDF of percent improvement over blob storage for a 128MB cache (left) and average percent improvement as the size of FaaS\$ varies (right). First, with just 128MB, FaaS\$ with pre-warm improves performance by 50% or more for about 35% of applications. FaaS\$ also has an average improvement of over 40% for a 128MB cache. Second, as the unload period gets smaller, FaaS\$’s pre-warm becomes more important to ensure frequently-accessed objects are available during the next application invocation. Third, improvement is correlated with reuse: we found that smaller objects tend to be reused more often, which resulted in greater performance improvements. Finally, we note that FaaS\$ is designed to support applications that run in production today (with object sizes of tens to hundreds of KB), but also for future applications that will access much larger objects (hundreds to thousands of MB) as shown in Section 6.3.

### 6.3 Comparing FaaS\$ to existing systems

**ML inference.** Figure 9 shows the latency for the AMBER Alert pipeline and single inference with AlexNet and SqueezeNet. First, the figure shows that FaaS\$ LH improves latency by 61%, 87%, and 55% faster than Vanilla for the AMBER Alert pipeline, AlexNet, and SqueezeNet, respectively. This demonstrates that avoiding remote storage accesses and using cache triggers can significantly improve FaaS\$ performance. Second, for the AMBER Alert pipeline, FaaS\$’s LH and RH are faster than using a Redis service, while FaaS\$ RH is equivalent to using a manually managed Redis VM (Pocket in Figure 9a). This is significant given the complexity of manually managing a Redis VM and the significant cost of using a Redis service (discussed below). FaaS\$ offers lower latency, while remaining transparent and relieving users of any configuration burden. Third, FaaS\$’s LM and RM exhibit similar

Notebook	Native IM	FaaS\$
Capacity planning	6.0s $\pm$ 0.2s	8.0s $\pm$ 0.3s
FaaS\$ characterization	40.4s $\pm$ 8.9s	68.0s $\pm$ 0.8s
No-Op	0.2ms $\pm$ 0ms	34.4ms $\pm$ 8.7ms

**Table 4: End-to-end latency running notebooks. FaaS\$ can run JupyterLess notebooks interactively.**

latencies, with the variability coming from the accesses to remote storage. This suggests using a multi-instance cache does not further penalize cache miss performance. Finally, FaaS\$ LH and RH perform well for both small (input images and class labels) and large objects (the models).

**JupyterLess.** Figure 10 shows the performance of summing a 350MB DataFrame column in a JupyterLess notebook. There are two Native setups: for In-Memory (IM) the DataFrame is pre-loaded in memory before summing, while for Remote Storage (RS) the latency of remote storage access is counted as part of the summation. FaaS\$ LH and RH improve performance compared to Vanilla by 62% and 29%, respectively. Compared to Native RS, FaaS\$’s LH and RH improve performance by 92% and 86%, respectively.

Table 4 shows the latency of running a capacity planning notebook (includes data collection and plotting), the FaaS\$ characterization of Section 2.1, and a No-Op notebook that logs a single message. Each run has a mix of data access scenarios for FaaS\$, since a local copy of the Jupyter state is saved per cell, and is read from a cachelet when executing the following cell. For the former two notebooks, the performance gap with Native IM comes from serializing plots and sending them to the notebook user interface. Thus, FaaS\$ can run JupyterLess notebooks interactively, and with near-native performance when the notebook is not trivial.

**Cost.** Figure 11 show the cost-per-hour comparison of running the applications end-to-end (lower is better). Cost is normalized to *Native*, and the y-axis is in log scale. We show cost for different application invocation IaTs between 10ms and 1 hour. For context, the application with the median IaT in our characterization (Section 2.1) has an average IaT of 20 minutes. For FaaS\$, we show the case where the application ran end-to-end with local hits. We do not show Redis service due to its large cost (FaaS\$ is 6 orders of magnitude cheaper).

The figure shows that FaaS\$ can provide huge cost savings. For applications with infrequent invocations (*e.g.*, once per hour), FaaS\$ is 99.999% cheaper than a Native VM. Per hour, FaaS\$ is cheaper than all baselines for all invocation intervals, except for the AMBER Alert pipeline at 10ms, where it is 33% more expensive than Native. For all other cases, FaaS\$ is 50% to 99.999% cheaper than caching and storage layers that require separate servers, such as Cloudburst and Pocket. From Section 2.1, 99.88% of applications have average IaT  $\geq$  10ms: in these cases, the cost of servers would almost always completely dominate the overall cost.

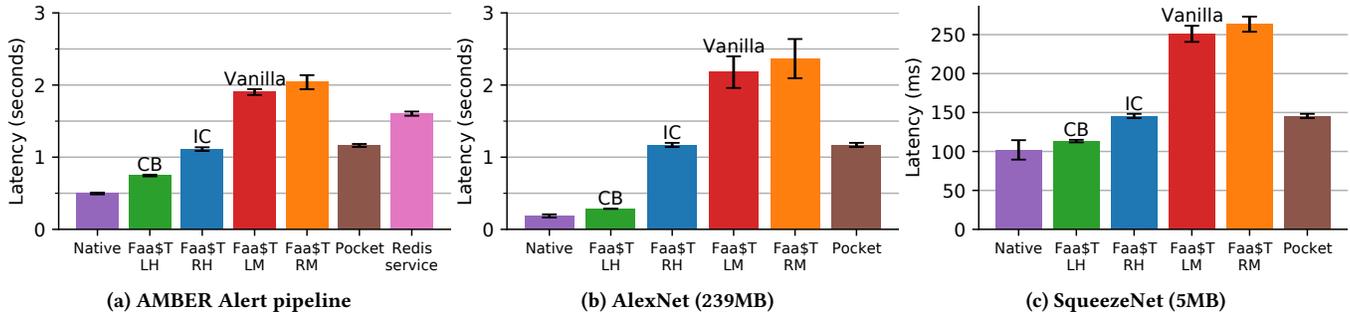


Figure 9: FaaS\$T versus existing systems for ML inference. FaaS\$T improves performance by accessing data in local or remote cache instances. (LH = Local Hit, RH = Remote Hit, LM = Local Miss, RM = Remote Miss, CB = Cloudburst, IC = InfiniCache.)

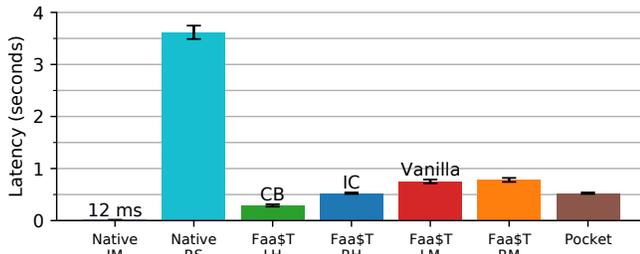


Figure 10: Summing a 350MB DataFrame column in a JupyterLess notebook. In Native IM, the DataFrame is already preloaded before summing, while Native RS fetches it from remote storage before summing. FaaS\$T improves performance by accessing data in local or remote cache instances.

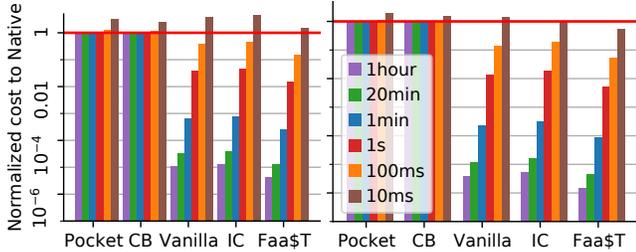


Figure 11: Cost normalized to Native VM for the AMBER Alert pipeline (left) and summing a 350MB DataFrame column in a JupyterLess notebook (right). Each bar represents an inter-invocation period. The y-axis is in log-scale (lower is better). FaaS\$T is 50% to 99.999% cheaper than the baselines.

**Discussion on comparison to Native.** Even when all accesses are served with local (LH) or remote hits (RH), FaaS\$T is slower than a Native VM with all data stored locally and there are no function invocation overheads. However, as we have shown in Figure 11, such a Native setup can be orders of magnitude more expensive than FaaS\$T, since we must keep all VMs running even when applications are idle. Moreover, the user is responsible for resource and data management. With FaaS\$T being a part of the FaaS runtime, users only pay for the time resources are consumed for both compute and caching. Moreover, local hits are on the order of hundreds of ms, which is close to interactive for many use cases.

Scenario	Heap growth succeeded?	Latency
No Mem Daemon	No	235.0ms $\pm$ 3.2ms
Mem Daemon, no scaling	Yes	678.6ms $\pm$ 64.8ms
Mem Daemon, scaling	Yes	502.9ms $\pm$ 28.5ms

Table 5: Latency of running a 350MB DataFrame summation in a JupyterLess notebook after growing heap memory, and whether the heap growth succeeded. We evaluate: (a) without the Memory Daemon, (b) with the Memory Daemon, but no cache size scaling, and (c) with Memory Daemon, and FaaS\$T scales to two instances. FaaS\$T does not compromise application functionality, and improves performance by scaling.

#### 6.4 Can FaaS\$T manage memory effectively?

We consider the JupyterLess notebook application that sums a 350MB column. After loading in the DataFrame and performing the summation, the application allocates an array that consumes 96% of the application's total memory. Then, the application again computes the summation of the 350MB column, which requires the reloading of the DataFrame that is evicted. We show three scenarios: (a) without the Memory Daemon to evict objects when the heap memory grows, (b) the Memory Daemon evicts objects, but there is no cache size scaling, and (c) the Memory Daemon evicts objects, and FaaS\$T scales to two instances.

Table 5 shows the performance of the second summation of the 350MB column, and whether the array allocation completed successfully. When the Memory Daemon does not trigger object eviction, the array allocation fails, but the summation matches the performance of a local hit. When the Memory Daemon triggers object eviction, but cannot scale, the array allocation succeeds, and the summation performance is a mix of local hits and misses. Finally, when the Memory Daemon triggers object eviction, and scales the cache size to improve the number of remote hits, the array allocation succeeds, and the summation performance is a mix of all four data access types. Since FaaS\$T opportunistically uses application memory, functionality is not compromised for applications that use large heap memory amounts.

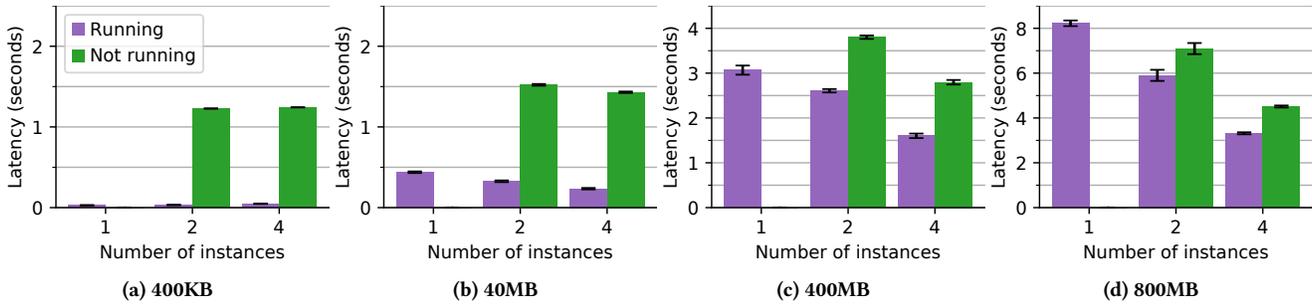


Figure 12: Latency of fetching an object from remote storage as the number of instances vary for increasingly large object sizes. If the instances are not loaded, they incur a cold-start; we only consider the running case for one instance. FaaS $\text{\$}$  determines whether to scale data loading across multiple instances to increase bandwidth.

## 6.5 Scaling as the object size varies

FaaS $\text{\$}$  can scale the number of instances based on object sizes. We consider four object sizes: 400KB, 40MB, 400MB, and 800MB. The amount of data downloaded by each FaaS $\text{\$}$  cachelet is evenly split between the available instances. For example, if there are two instances and the object size is 400KB, each one downloads 200KB. The data is then processed at a single instance. For each object size, we show two cases: (a) when all application instances are running, and (b) when additional instances (more than one) must be loaded in order to fetch the objects (thus incurring a cold-start latency).

Figure 12 shows the time to download the increasingly large object sizes from remote storage as the number of instances varies. For small objects (400KB), the download latency is low enough that using more than one instance degrades performance, especially if the instances need to be loaded. For the 40MB object, partitioning the download across multiple instances is beneficial if the instances are already loaded. It is 47% faster to download the object with four cachelets than with one. For the 400MB object, it is 9% faster to download the object with four instances than one if the instances are not already loaded, and 47% faster to download the objects with four instances than one if the instances are already loaded. For the 800MB object, it is beneficial to use four cachelets to load the object in parallel, even if the instances need to be loaded. If the instances are already loaded, it is 60% faster to download with four cachelets compared to one. If the instances need to be loaded, it is 44% faster.

## 6.6 Consistency versus performance

FaaS $\text{\$}$  allows users to trade off consistency for performance. We evaluate the write/read settings from Table 3 using a JupyterLess notebook in which five cachelets share the state for counting up to 1K. Application instances add to the counter in round-robin fashion, and we expect the counter at the end to have a value of 1K. We measure inconsistencies as the absolute difference between the final counter value and 1K. This is a critical primitive in multiplayer games [17].

Write target	Write mode	Read target	E2E lat (s)	Per-req. lat (ms)	# inconsist.
Storage	Sync	Storage	74.4 ± 0.2	80.0 ± 29.5	0 ± 0
Owner	Sync	Owner	42.6 ± 0.2	41.6 ± 12.8	0 ± 0
Owner	Async	Owner	39.7 ± 0.6	38.5 ± 11.4	1.3 ± 1.9
Owner	Sync	Local	32.7 ± 0.3	32.7 ± 9.0	800 ± 0
Local	Sync	Local	31.1 ± 0.3	30.3 ± 13.1	800 ± 0
Local	Aync	Local	31.0 ± 0.3	30.3 ± 13.2	800 ± 0

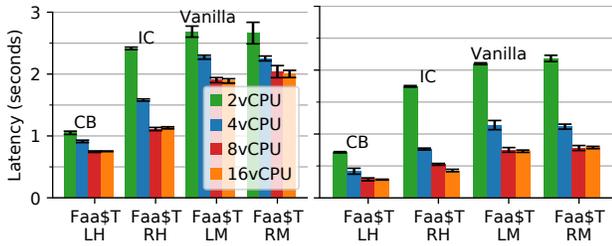
Table 6: Latency (end-to-end and per-request) and number of inconsistencies for different settings for a JupyterLess notebook counting to 1K with five instances sharing state. Inconsistencies are the absolute difference between the final counter value and 1K. Performance increases as consistency and fault tolerance decrease.

Table 6 shows the end-to-end latency, per-request latency, and number of inconsistencies for all five settings. As expected, latency drops as we relax consistency requirements. For example, writing to the local cache and reading from the local cache is fastest, but provides no consistency (800 inconsistencies recorded) and the lowest fault tolerance. Writing to the local cache asynchronously and reading from the local cache is equivalent to Cloudburst’s performance. Cloudburst would exhibit better consistency due to its lattice datatypes, but requires support from the datastore. Latency varies the most when writing and reading from remote storage, and the least when writing and reading from the local cache.

## 6.7 Sensitivity to instance size

Finally, we evaluate the sensitivity of running applications with FaaS $\text{\$}$  as the instance size varies. We run the AMBER Alert pipeline and sum a 350MB DataFrame column in a JupyterLess notebook. Instances scale linearly in terms of memory and network bandwidth as the number of vCPUs increases. For example, the 2vCPU instance has 8GiB of memory and 1Gbps of network bandwidth (Standard\_DS2\_v2), and the 4vCPU instance has 16GiB of memory and 2Gbps of network bandwidth (Standard\_DS3\_v2) [7].

Figures 13a and 13b show that, as instances increase in size, latency decreases for both applications. Larger instances have higher network bandwidth, which is beneficial for data



(a) AMBER Alert pipeline (b) 350MB DataFrame Notebook

**Figure 13: Latency as instance size varies for (a) the AMBER Alert pipeline, and (b) summing a 350MB DataFrame column in a JupyterLess notebook. Instances memory and network bandwidth scale linearly as the number of cores increases. FaaS $T$  benefits from higher bandwidth instances.**

accesses to remote storage and between instances. Data accesses of these two applications saturate the bandwidth with the 8vCPU instance. Thus, although the 16vCPU instance is the highest bandwidth instance size, the performance remains the same as the 8vCPU. Some cloud providers offer instances as small as 2vCPU with up to 10Gbps bandwidth, allowing FaaS $T$  to have high performance even on small instances suitable for FaaS.

## 7 DISCUSSION

**Stateful serverless environments.** Stateful FaaS platforms (e.g., Azure Durable Functions [12]) are designed to manage and persist coordination state about the execution. While this diminishes their need for remote storage, they still require remote storage for input and output data, snapshots, persisting progress, and communication, which can be cached and managed by FaaS $T$ . Hence, FaaS $T$ 's design would likely stay the same even with the emergence of stateful FaaS platforms as it (a) enables transparent communication between cachelets, which is not offered by existing stateful platforms, and (b) has features like pre-warming and bandwidth scaling that are necessary even with alternative means of communication.

**SSD versus in-memory caching.** We made FaaS $T$  an in-memory cache since we found an SSD-only cache was not performant enough to enable interactive applications (especially for small objects). However, it makes sense to explore a hierarchical/tiered approach where FaaS $T$  optimizes for cost and performance based on object access patterns.

## 8 RELATED WORK

**Ephemeral serverless storage.** In Section 2.2 we describe the limitations of several existing storage and data caching solutions for FaaS [34, 45, 51, 53, 56]. Unlike these systems, FaaS $T$  does not require external resources beyond what is provided to the function, is transparent to applications, and can scale as the data size and access patterns vary.

OFC [39] is the closest work to FaaS $T$ . It transparently caches objects using RAMCloud [41] and leverages machine learning to dynamically size the cache. Unlike OFC, FaaS $T$

pre-warms objects when an application is loaded, supports large (> 10MB) object caching and scales to reduce their data transfer latency from remote storage, and only needs to keep one copy of data in shared memory (compared to OFC that requires a copy in the worker and in RAMCloud). FaaS $T$  also incurs lower decision overheads and is easier to manage by not requiring machine learning for its decision-making.

**Serverless frameworks.** Several frameworks have recently emerged enabling users to run, for example, linear algebra [30], video encoding [20], video analytics pipelines [4], ML training [14], and general burst-parallel applications [19] on up to thousands of serverless functions. These, and their applications, would benefit from managing and transferring intermediate data between serverless functions using FaaS $T$ . Since FaaS $T$  is transparent to applications, little to no changes would be needed to interact with FaaS $T$ .

**Improving serverless performance.** Many approaches have been proposed to reduce serverless function execution times: making containers more lightweight [1, 40], snapshotting techniques [13, 18, 55], and reducing the number of cold-starts [21, 48]. Shredder [60] focuses on how to provide isolation for multi-tenancy. Lambada [38] improves performance for applications with exchange operators. Lambdata [54] allows users to expose their data read and write intents to co-locate functions working on the same data. FaaS $T$  is orthogonal and compatible with these optimizations.

**Consistency and fault tolerance.** Consistency and fault tolerance protocols have been heavily studied. Recent work has explored how to enable both of these for serverless applications using the KVS backend [51, 53], fault tolerance shim layers [52], and novel data structures to support exactly-one semantics [59]. FaaS $T$  transparently supports different consistency and fault tolerance settings directly in the functions runtime, and can be extended to support future protocols.

## 9 CONCLUSION

FaaS $T$  is a transparent caching layer for serverless applications. We motivated its design with a characterization of production applications. We tie FaaS $T$  to the application, scale based on compute demands and data access patterns, and provide data consistency that can be set per application. We implemented it for Azure Functions and showed that compared to existing systems, FaaS $T$  is on average 57% faster and 99.99% cheaper when running two challenging applications. We released sanitized logs from our characterization at [6].

## Acknowledgements

We thank our shepherd, Orran Krieger, the anonymous reviewers, Daniel Berger, Phil Bernstein, Sameh Elnikety, Jose Faleiro, Jeff Hollan, and Stanko Novakovic for their many helpful comments and suggestions. This work was supported by the Stanford Platform Lab and its industrial affiliates.

## REFERENCES

- [1] Alexandru Agache, Marc Brooker, Alexandra Iordache, Anthony Liguori, Rolf Neugebauer, Phil Piwonka, and Diana-Maria Popa. 2020. Firecracker: Lightweight Virtualization for Serverless Applications. In *NSDI*.
- [2] Amazon S3 2021. Amazon S3. <https://aws.amazon.com/s3/>.
- [3] Amazon Web Services. 2021. AWS SDK for Python (Boto3). <https://aws.amazon.com/sdk-for-python/>.
- [4] Lixiang Ao, Liz Izhikevich, Geoffrey M. Voelker, and George Porter. 2018. Sprocket: A Serverless Video Processing Framework. In *SoCC*.
- [5] Gabriel Aumala, Edwin Boza, Luis Ortiz-Avilés, Gustavo Totoy, and Cristina Abad. 2019. Beyond Load Balancing: Package-Aware Scheduling for Serverless Platforms. In *CCGRID*.
- [6] Microsoft Azure and Microsoft Research. 2021. Azure Functions Traces. <https://github.com/Azure/AzurePublicDataset>.
- [7] Azure DSv2-series 2021. Azure DSv2-series. <https://docs.microsoft.com/en-us/azure/virtual-machines/dv2-dsv2-series#dsv2-series>.
- [8] Azure Function Queues 2021. Azure Function Queues. <https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue>.
- [9] Nathan Beckmann, Haoxian Chen, and Asaf Cidon. 2018. LHD: Improving Cache Hit Rate by Maximizing Hit Density. In *NSDI*.
- [10] N. Beckmann and D. Sanchez. 2017. Maximizing Cache Performance Under Uncertainty. In *HPCA*.
- [11] Daniel S. Berger, Ramesh K. Sitaraman, and Mor Harchol-Balter. 2017. AdaptSize: Orchestrating the Hot Object Memory Cache in a Content Delivery Network. In *NSDI*.
- [12] Sebastian Burckhardt, Chris Gillum, David Justo, Konstantinos Kallas, Connor McMahon, and Christopher S. Meiklejohn. 2021. Serverless Workflows with Durable Functions and Netherite. *arXiv* (2021).
- [13] James Cadden, Thomas Unger, Yara Awad, Han Dong, Orran Krieger, and Jonathan Appavoo. 2020. SEUSS: Skip Redundant Paths to Make Serverless Fast. In *EuroSys*.
- [14] Joao Carreira, Pedro Fonseca, Alexey Tumanov, Andrew Zhang, and Randy Katz. 2019. Cirrus: A serverless framework for end-to-end ML workflows. In *SoCC*.
- [15] Gohar Irfan Chaudhry. 2021. Caching of function bindings (input/output) for faster I/O. <https://github.com/Azure/azure-functions-host/issues/7310>.
- [16] Gohar Irfan Chaudhry. 2021. Shared memory data transfer between Functions Host and out-of-proc workers. <https://github.com/Azure/azure-functions-host/pull/6836>.
- [17] Jesse Donkervliet, Animesh Trivedi, and Alexandru Iosup. 2020. Towards Supporting Millions of Users in Modifiable Virtual Environments by Redesigning Minecraft-Like Games as Serverless Systems. In *HotCloud*.
- [18] Dong Du, Tianyi Yu, Yubin Xia, Binyu Zang, Guanglu Yan, Chenggang Qin, Qixuan Wu, and Haibo Chen. 2020. Catalyzer: Sub-Millisecond Startup for Serverless Computing with Initialization-Less Booting. In *ASPLOS*.
- [19] Sadjad Fouladi, Francisco Romero, Dan Iter, Qian Li, Shuvo Chatterjee, Christos Kozyrakis, Matei Zaharia, and Keith Winstein. 2019. From Laptop to Lambda: Outsourcing Everyday Jobs to Thousands of Transient Functional Containers. In *USENIX ATC*.
- [20] Sadjad Fouladi, Riad S. Wahby, Brennan Shacklett, Karthikeyan Vasuki Balasubramaniam, William Zeng, Rahul Bhalerao, Anirudh Sivaraman, George Porter, and Keith Winstein. 2017. Encoding, Fast and Slow: Low-Latency Video Processing Using Thousands of Tiny Threads. In *NSDI*.
- [21] Alexander Fuerst and Prateek Sharma. 2021. FaasCache: Keeping Serverless Computing Alive With Greedy-Dual Caching. In *ASPLOS*.
- [22] U. Gupta, S. Hsia, V. Saraph, X. Wang, B. Reagen, G. Wei, H. S. Lee, D. Brooks, and C. Wu. 2020. DeepRecSys: A System for Optimizing End-To-End At-Scale Neural Recommendation Inference. In *ISCA*.
- [23] Udit Gupta, Carole-Jean Wu, Xiaodong Wang, Maxim Naumov, Brandon Reagen, David Brooks, Bradford Cottel, Kim Hazelwood, Mark Hempstead, Bill Jia, et al. 2020. The Architectural Implications of Facebook’s DNN-Based Personalized Recommendation. In *HPCA*.
- [24] K. Hazelwood, S. Bird, D. Brooks, S. Chintala, U. Diril, D. Dzhulgakov, M. Fawzy, B. Jia, Y. Jia, A. Kalro, J. Law, K. Lee, J. Lu, P. Noordhuis, M. Smelyanskiy, L. Xiong, and X. Wang. 2018. Applied Machine Learning at Facebook: A Datacenter Infrastructure Perspective. In *HPCA*.
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. In *CVPR*.
- [26] Joseph M Hellerstein, Jose Faleiro, Joseph E Gonzalez, Johann Schleier-Smith, Vikram Sreekanti, Alexey Tumanov, and Chenggang Wu. 2018. Serverless Computing: One Step Forward, Two Steps Back. *arXiv preprint arXiv:1812.03651* (2018).
- [27] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *arXiv* (2017).
- [28] Forrest N. Iandola, Matthew W. Moskewicz, Khalid Ashraf, Song Han, William J. Dally, and Kurt Keutzer. 2016. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <1MB model size. *CoRR* abs/1602.07360 (2016). [arXiv:1602.07360](http://arxiv.org/abs/1602.07360) <http://arxiv.org/abs/1602.07360>
- [29] Fei Jiang, Yong Jiang, Hui Zhi, Yi Dong, Hao Li, Sufeng Ma, Yilong Wang, Qi ang Dong, Haipeng Shen, and Yongjun Wang. 2017. Artificial intelligence in healthcare: past, present and future. *Stroke and Vascular Neurology* 2, 4 (2017), 230–243. <https://doi.org/10.1136/svn-2017-000101> [arXiv:https://svn.bmj.com/content/2/4/230.full.pdf](https://svn.bmj.com/content/2/4/230.full.pdf)
- [30] Eric Jonas, Qifan Pu, Shivaram Venkataraman, Ion Stoica, and Benjamin Recht. 2017. Occupy the Cloud: Distributed Computing for the 99%. In *SoCC*.
- [31] Eric Jonas, Johann Schleier-Smith, Vikram Sreekanti, Chia-Che Tsai, Anurag Khandelwal, Qifan Pu, Vaishaal Shankar, Joao Carreira, Karl Krauth, Neeraja Yadwadkar, et al. 2019. Cloud programming simplified: A Berkeley view on serverless computing. *arXiv* (2019).
- [32] Jupyter 2021. Jupyter. <https://jupyter.org/>.
- [33] David Karger, Eric Lehman, Tom Leighton, Rina Panigrahy, Matthew Levine, and Daniel Lewin. 1997. Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the world wide web. In *STOC*.
- [34] Ana Klimovic, Yawen Wang, Patrick Stuedi, Animesh Trivedi, Jonas Pfefferle, and Christos Kozyrakis. 2018. Pocket: Elastic ephemeral storage for serverless analytics. In *OSDI*.
- [35] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*.
- [36] Matt McFarland. 2017. Futuristic cop cars may identify suspects. <https://money.cnn.com/2017/10/19/technology/future/police-ai-dashcam/index.html>.
- [37] Microsoft Azure Blob Storage 2021. Microsoft Azure Blob Storage. <https://azure.microsoft.com/en-us/services/storage/blobs/>.
- [38] Ingo Müller, Renato Marroquín, and Gustavo Alonso. 2020. Lambda: Interactive Data Analytics on Cold Data Using Serverless Cloud Infrastructure. In *SIGMOD*.
- [39] Djob Mvondo, Mathieu Bacou, Kevin Nguetchouang, Stephane Pouget, Josiane Kouam, Renaud Lachaize, Jinho Hwang, Tim Wood, Daniel Hagimont, Noel De Palma, Bernabe Batchakui, and Alain Tchana. 2021. OFC: An Opportunistic Caching System for FaaS Platforms. In *EuroSys*.
- [40] Edward Oakes, Leon Yang, Dennis Zhou, Kevin Houck, Tyler Harter, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau. 2018. SOCK:

- Rapid Task Provisioning with Serverless-Optimized Containers. In *USENIX ATC*.
- [41] John Ousterhout, Arjun Gopalan, Ashish Gupta, Ankita Kejriwal, Collin Lee, Behnam Montazeri, Diego Ongaro, Seo Jin Park, Henry Qin, Mendel Rosenblum, Stephen Rumble, Ryan Stutsman, and Stephen Yang. 2015. The RAMCloud Storage System. *ACM Trans. Comput. Syst.* (2015).
- [42] Protocol Buffers 2021. Protocol Buffers. <https://developers.google.com/protocol-buffers/>.
- [43] Fotis Psallidas, Yiwen Zhu, Bojan Karlas, Matteo Interlandi, Avriela Floratou, Konstantinos Karanasos, Wentao Wu, Ce Zhang, Subru Krishnan, Carlo Curino, and Markus Weimer. 2019. Data Science through the looking glass and what we found there. arXiv:1912.09536 [cs.LG]
- [44] Qifan Pu, Haoyuan Li, Matei Zaharia, Ali Ghodsi, and Ion Stoica. 2016. FairRide: Near-Optimal, Fair Cache Sharing. In *NSDI*.
- [45] Qifan Pu, Shivaram Venkataraman, and Ion Stoica. 2019. Shuffling, Fast and Slow: Scalable Analytics on Serverless Infrastructure. In *NSDI*.
- [46] V. J. Reddi, C. Cheng, D. Kanter, P. Mattson, G. Schmuelling, C. Wu, B. Anderson, M. Breughe, M. Charlebois, W. Chou, R. Chukka, C. Coleman, S. Davis, P. Deng, G. Damos, J. Duke, D. Fick, J. S. Gardner, I. Hubara, S. Idgunji, T. B. Jablin, J. Jiao, T. S. John, P. Kanwar, D. Lee, J. Liao, A. Lokhmotov, F. Massa, P. Meng, P. Micikevicius, C. Osborne, G. Pekhimenko, A. T. R. Rajan, D. Sequeira, A. Sirasao, F. Sun, H. Tang, M. Thomson, F. Wei, E. Wu, L. Xu, K. Yamada, B. Yu, G. Yuan, A. Zhong, P. Zhang, and Y. Zhou. 2020. MLPerf Inference Benchmark. In *ISCA*.
- [47] Redis 2021. Redis. <https://redis.io/>.
- [48] Mohammad Shahradd, Rodrigo Fonseca, Inigo Goiri, Gohar Irfan, Paul Batum, Jason Cooke, Eduardo Laureano, Colby Tresness, Mark Rusinovich, and Ricardo Bianchini. 2020. Serverless in the Wild: Characterizing and Optimizing the Serverless Workload at a Large Cloud Provider. In *USENIX ATC*.
- [49] Vaishaal Shankar, Karl Krauth, Kailas Vodrahalli, Qifan Pu, Benjamin Recht, Ion Stoica, Jonathan Ragan-Kelley, Eric Jonas, and Shivaram Venkataraman. 2020. Serverless Linear Algebra. In *SoCC*.
- [50] Mikhail Shilkov. 2021. Comparison of Cold Starts in Serverless Functions across AWS, Azure, and GCP. <https://mikhail.io/serverless/coldstarts/big3/>.
- [51] Simon Shillaker and Peter Pietzuch. 2020. Faasm: Lightweight Isolation for Efficient Stateful Serverless Computing. In *USENIX ATC*.
- [52] Vikram Sreekanti, Chenggang Wu, Saurav Chhatrapati, Joseph E. Gonzalez, Joseph M. Hellerstein, and Jose M. Faleiro. 2020. A Fault-Tolerance Shim for Serverless Computing. In *EuroSys*.
- [53] Vikram Sreekanti, Chenggang Wu, Xiayue Charles Lin, Johann Schleier-Smith, Joseph E. Gonzalez, Joseph M. Hellerstein, and Alexey Tumanov. 2020. Cloudburst: Stateful Functions-as-a-Service.
- [54] Y. Tang and J. Yang. 2020. Lambda: Optimizing Serverless Computing by Making Data Intents Explicit. In *CLOUD*.
- [55] Dmitrii Ustiugov, Plamen Petrov, Marios Kogias, Edouard Bugnion, and Boris Grot. 2021. Benchmarking, Analysis, and Optimization of Serverless Function Snapshots. In *ASPLOS*.
- [56] Ao Wang, Jingyuan Zhang, Xiaolong Ma, Ali Anwar, Lukas Rupprecht, Dimitrios Skourtis, Vasily Tarasov, Feng Yan, and Yue Cheng. 2020. InfiniCache: Exploiting Ephemeral Serverless Functions to Build a Cost-Effective Memory Cache. In *FAST*.
- [57] C. Wu, J. Faleiro, Y. Lin, and J. Hellerstein. 2018. Anna: A KVS for Any Scale. In *ICDE*.
- [58] Neeraja J. Yadwadkar, Francisco Romero, Qian Li, and Christos Kozyrakis. 2019. A Case for Managed and Model-Less Inference Serving. In *HotOS*.
- [59] Haoran Zhang, Adney Cardoza, Peter Baile Chen, Sebastian Angel, and Vincent Liu. 2020. Fault-tolerant and transactional stateful serverless workflows. In *OSDI*.
- [60] Tian Zhang, Dong Xie, Feifei Li, and Ryan Stutsman. 2019. Narrowing the Gap Between Serverless and Its State with Storage Functions. In *SoCC*.
- [61] Yunqi Zhang, George Prekas, Giovanni Matteo Fumarola, Marcus Fountoura, Inigo Goiri, and Ricardo Bianchini. 2016. History-Based Harvesting of Spare Cycles and Storage in Large-Scale Datacenters. In *OSDI*.