

SPES: Towards Optimizing Performance-Resource Trade-Off for Serverless Functions

Cheryl Lee*, Zhouruixing Zhu[†], Tianyi Yang*, Yintong Huo*, Yuxin Su[‡], Pinjia He[†], and Michael R. Lyu*

*The Chinese University of Hong Kong, Hong Kong, China.

Email: cherylle@link.cuhk.edu.hk, {tyyang, ythuo, lyu}@cse.cuhk.edu.hk

[†]The Chinese University of Hong Kong, Shenzhen, Shenzhen, China.

Email: zhouruixingzhu@link.cuhk.edu.cn, hepinjia@cuhk.edu.cn

[‡]School of Software Engineering, Sun Yat-sen University, Zhuhai, China. Email: suyx35@mail.sysu.edu.cn

Abstract—As an emerging cloud computing deployment paradigm, serverless computing is gaining traction due to its efficiency and ability to harness on-demand cloud resources. However, a significant hurdle remains in the form of the *cold start* problem, causing latency when launching new function instances from scratch. Existing solutions tend to use oversimplistic strategies for function pre-loading/unloading without full invocation pattern exploitation, rendering unsatisfactory optimization of the trade-off between cold start latency and resource waste. To bridge this gap, we propose SPES, the first differentiated scheduler for runtime cold start mitigation by optimizing serverless function provision. Our insight is that the common architecture of serverless systems prompts the concentration of certain invocation patterns, leading to predictable invocation behaviors. This allows us to categorize functions and pre-load/unload proper function instances with finer-grained strategies based on accurate invocation prediction. Experiments demonstrate the success of SPES in optimizing serverless function provision on both sides: reducing the 75th-percentile cold start rates by 49.77% and the wasted memory time by 56.43%, compared to the state-of-the-art. By mitigating the cold start issue, SPES is a promising advancement in facilitating cloud services deployed on serverless architectures.

Index Terms—Cloud computing, serverless computing, cold start, function categorization

I. INTRODUCTION

Function-as-a-service (FaaS), the most prominent implementation pattern of serverless computing, has extensively simplified developers' access to cloud resources. Major cloud vendors such as AWS Lambda, Google Cloud Functions, and Azure Functions have supported FaaS-based web services [1], [2], machine learning [3], [4], and other cloud applications [5], [6]. FaaS shifts the burden of infrastructure management from developers to cloud vendors, allowing developers to focus on their application functionality [7]. The statelessness and event-driven architecture of FaaS facilitate seamless updates and flexible application deployment. Plus, FaaS applies a pay-as-you-go billing model and dynamically allocates resources based on demand [8], appealing to users via cost-saving benefits. Figure 1 depicts a serverless application example.

Despite their benefits, cloud users on serverless platforms suffer from the infamous cold-start problem, which causes them to trade off performance latency and memory costs. Serverless functions are spawned on instances and typically have very short execution duration (orders of milliseconds




Fig. 1. An weather inquiry website based on a serverless web application.

to seconds) [9]. However, booting a function from scratch (*i.e.*, *cold start*) incurs expensive latency in preparing the execution environment [10], compared to running a *warm* function whose instance is already loaded in the memory. Cold-start latency can account for 80% of the total response latency [11], leading to user dissatisfaction and eventually causing user churn and economic loss. The pay-per-use model of FaaS also motivates cloud vendors to avoid unbillable set-up latency [12]. Meanwhile, keeping all functions warm is infeasible due to two reasons. First, most functions are invoked infrequently [13]. It is unacceptable to keep idle (also unbillable) and infrequently used function instances in memory, as it leads to unnecessary costs and wastes valuable resources. Second, because of the infrequent invocation and short execution, FaaS providers colocate thousands of function instances on a single server to achieve full utilization [14]. Loading all functions can occupy hundreds of GBs of memory, far beyond a server's capacity [14]. Thus, reducing cold start latency while minimizing memory waste has been a key challenge [15], [3], [16].

Existing solutions generally fall into two categories [17], [18]: accelerating the setup [19], [20], [21] or reducing cold start occurrences [22], [13], [23]. This paper focuses on the reduction of runtime cold start occurrence without renovating the underlying system-layer infrastructure. Previous such efforts [22], [13], [23] either keep instances loaded for a fixed period post-execution or manage instances based on rudimentary inter-invocation intervals. However, they overlook the intrinsic influence of triggering events on invocation patterns, thus falling short in addressing infrequently invoked functions. Hence, we intend to develop a more effective, non-intrusive method that can load functions right before incoming invocations and recycle idle instances with no potential near-

future invocations. The key is to predict invocations accurately.

We identify four challenges in developing such a method:

- *Efficiency*: A FaaS platform can receive thousands of function invocations every minute from millions of potential functions [24]. Most functions have a very short execution time, so the method must decide the provision within a limited and unbillable time.
- *Scalability*: The method should be highly elastic, quickly scaling up and down to meet invocation fluctuations promptly, as the requests of FaaS functions are usually bursty and dynamic [2].
- *Imbalance*: The distribution of function invocations can be highly imbalanced, and most functions are rarely invoked [13]. This poses a challenge to approaches requiring sufficient training data.
- *Evolution*: The lightweight nature of FaaS makes applications evolve faster than traditional software. Both short-term factors (e.g., user activity variations) and long-term factors (e.g., software updates) can cause concept shifts in function invocations, thereby hindering predictive models.

Our insight: *the event-driven architecture of serverless systems enables identifying common invocation patterns, allowing for efficient function categorization and provision with bespoke strategies.* This insight is grounded in both domain expertise and data analysis. Serverless functions are triggered by business-specific events and often fall into limited types, laying a foundation for invocation prediction. For example, the invocations of a scheduled data processing function are easy to predict, and so is the pre-loading. While provision is trigger-agnostic, invocation patterns can be modeled from historical invocations associated with triggers. Moreover, functions within an application often display behavior convergence. For instance, functions within a multi-stage processing workload are typically invoked in turn. Our empirical analysis also reveals such consistent behavior patterns. Besides, some functions with irregular invocations may exhibit temporal locality, making it feasible to keep them loaded temporarily, reducing cold starts with minimal memory overhead. Thus, our study shows that functions can present limited predictable patterns, allowing for differentiated strategic pre-loading to minimize cold starts with minimal memory impact.

Hence, we propose **SPES**, a differentiated Scheduler for Provisioning runtime E Serverless functions in order to optimize the trade-off between performance latency and resource cost. After extensive data analysis, we summarize five typical invocation patterns and design three extra policies for functions experiencing less predictive invocations. We also propose a simple but effective metric, *i.e.* co-occurrence rate, to measure the correlation degree between functions. Highly correlated functions can be predictive indicators of each other. In addition, we design adaptive strategies to handle the concept shifts as FaaS evolves. Hence, SPES can predict the next invocation accurately and provision the function in advance with fine-grained strategies. As a result, our rule-based SPES addresses all the above challenges. Its tailored strategies can handle diverse functions with remarkable scalability and adaptivity,

introducing a small computation overhead.

We evaluate SPES on the most widely used industrial dataset released by Microsoft Azure [13]. Experiments demonstrate the superiority of SPES, which achieves both-side improvement: reducing the 75th-percentile cold start rate by 49.77%–89.20% and the wasted memory time by 13.51%–64.58% compared to baselines, and also reducing the proportion of “always-cold” functions significantly. Extensive ablation studies further confirm the effectiveness of our designs of inter-function connection and adaptive strategies.

In summary, the main contributions of this paper are:

- We propose SPES, the first differentiated method to mitigate the cold start issue by optimizing the runtime provision of serverless functions based on invocation prediction.
- We summarize typical invocation patterns and categorize functions accordingly as the basis of tailored provision strategies. We also create an effective “co-occurrence” metric to gauge inter-function correlations. Our code is available at <https://github.com/BEbillionaireUSD/SPES>.
- We conduct extensive experiments to demonstrate the significant improvement of SPES in the cold start and wasted memory reduction, as well as the contribution of designs for adaptivity and unseen functions.

II. BACKGROUND

A. Serverless Computing

Serverless computing represents an emerging cloud programming paradigm where cloud providers fully manage the underlying infrastructure and allocate resources dynamically, enabling developers to concentrate solely on their application’s core logic. In FaaS, developers implement services as stateless workloads, *i.e.* functions, designed to respond to an individual event using pre-defined rules, called the *trigger* [25]. Developers only pay for the actual computation resources on a per-execution basis, making it cost-efficient and scalable [26]. Besides, a cloud service-based *application* is usually broken down into separate, independent functions in practical implementation, where these functions are sometimes explicitly chained together [27]. The adoption of serverless computing has witnessed rapid growth, especially in web services, machine learning training [28], etc. Notwithstanding its advantages, serverless computing poses new challenges, such as ensuring that a service adheres to quality of service (QoS) demands related to response time or tail latency [29].

B. Cold Start Challenge

Figure 2 displays a serverless function’s lifecycle. A FaaS provider first downloads the code of users and initiates the execution environment (or sandbox) in the memory on cluster machines, known as a *cold start*. A cold start involves retrieving code from storage, container (or VM) initialization, loading code into memory, and executing the function’s handler. In contrast, a *warm start* jumps directly to execution, resulting in a faster response. Upon serving a request, its environment remains idle for a while without other invocations before the orchestration system decommissions it [30]. Naturally,

the more reused environments, the less cold starts. There exists a trade-off in keeping the instance alive: saving start-up resources and speeding up subsequent requests but incurring *idle time* costs. Such costs can be quantified by *wasted memory time (WMT)*, i.e., the time when the image of a function is kept in memory, but the function is not actually invoked. WMT is an important metric to measure resource waste in practice [13].




Fig. 2. A serverless function’s lifecycle.

The cold start problem is critical because it significantly increases runtime latency, yet serverless providers should meet the QoS requirements and maintain an “always-ready” illusion to users. On the one hand, cold starts can dominate the overall execution time [21], [15], [31]. [19] shows that the “Execution/Overall” latency ratio of most tested functions in gVisor can not even achieve 30%. On the other hand, the cold start happens very frequently [32], [18], [33]. Hence, the additional latency incurred by cold starts is exceptionally unbearable. Considering that memory is limited and expensive, mitigating the cold start challenge aims to either load/unload proper function instances or speed up cold start initiation.

III. PRELIMINARY EMPIRICAL ANALYSIS

A. Challenges of Serverless Function Provision

1) *Efficiency requirement*: Serverless provision must make immediate decisions for massive invocations, which hinders the potential of applying complex models. This requirement stems from function massiveness and latency sensitivity. Large-scale FaaS platforms launch a burst-parallel swarm and allow for thousands of invocations every minute [34], all waiting for the provision decision. Moreover, serverless functions are usually short-lived, so a provisioning method must handle invocations and update its strategy quickly during the unbillable scheduling time. As studied in [35], it is better to schedule invocations for execution as soon as they enter the system. A fast and simple provision method is thereby much more desired.

2) *Scalability under invocation spikes*: The provisioning method is also expected to scale automatically to a very high load. FaaS invocations are extremely bursty and dynamic [33]. An ideal FaaS platform should handle the invocation burst, so the provision decision-maker must be capable of rapid and elastic scaling to respond to sudden invocation spikes seamlessly and scale down during periods of low demand. For example, during a holiday shopping period, an e-commerce site running on a FaaS architecture may experience a 10x increase in traffic and function invocations compared to a normal workday. Hence, we require a highly scalable solution.

3) *Imbalance in invocation distribution*: Serverless functions exhibit heterogeneity in functionality, manifesting distinct implementations through a diversity of libraries and run-

times [29]. Moreover, the functions are triggered in response to various events, such as HTTP requests, data updates, and the receipt of notifications [36]. The various real-world demands make the frequencies of function invocations dramatically different. Figure 3 shows the distribution of invocations from an industrial dataset released by Microsoft Azure [13], collected for 14 days. We can observe that the invocations are highly non-uniform, with most functions being rarely invoked. This imbalance poses a challenge to learning-based approaches that require lots of records for training because many functions may not have invocation records in the data collection.

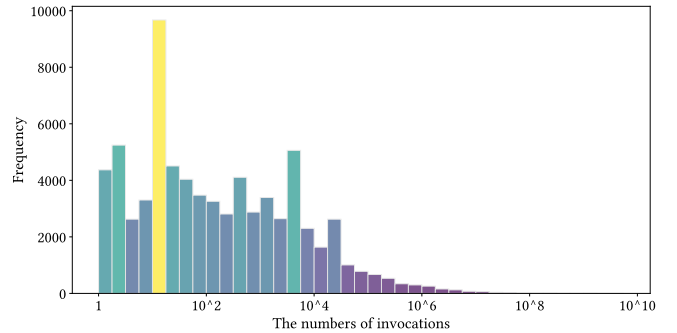


Fig. 3. The distribution of function invocations. The x-axis represents the range of function invocation counts, and the y-axis indicates the number of functions falling into the corresponding ranges.

4) *Evolution in invocation behavior*: Invocation evolution can degrade the performance of provisioning strategies because they rely on proactively pre-loading to-be-invoked functions to reduce cold starts fundamentally, whose predictive modeling can result in poor performance with invocation evolution. The situation is even worse for learning-based approaches, which may require frequent, time-consuming re-training. Many factors can cause such shifts, e.g. seasonal traffic, data volume changes, geographic expansion, etc. Considering the e-commerce site mentioned earlier, it can experience short-term spikes during holiday seasons and evolve to long-term increases when expanding to overseas markets. Changes in business logic are eventually propagated to the function development and utilization, resulting in observable shifts in invocations. In the scope of a single serverless function, a function can be re-bound to a new trigger or new calling chains, leading to noticeable concept shifts. We also observe the short-term evolution of invocations in the mentioned dataset, as shown in Figure 4, which plots the invocations of three typical functions over time. This fast-evolving characteristic of FaaS requires an adaptive provision method capable of handling concept shifts.

B. Observations and Our Insight

1) *Invocation pattern and triggers*: Serverless systems operate on an event-driven architecture. Generally, a trigger defines how a function is invoked by causing a function to run, so function invocation patterns are decided by pre-defined trigger types, such as HTTP requests, queued messages, and scheduled timers, resulting in predictive behaviors. Two observations from the mentioned dataset support this argument:




Fig. 4. The function invocations can experience distinct concept shifts that may degrade the provision performance. Different colors distinguish changes over time in the function invocation patterns.

First, most functions correspond to only one trigger type. The proportion of triggers bound with functions is shown in Figure 5, where “combination” means more than one (at most three) type of triggers bound with one function. Since only 2.6% functions have more than one trigger type, we can link each function to one pattern without concerning too much about the derived combination of patterns.

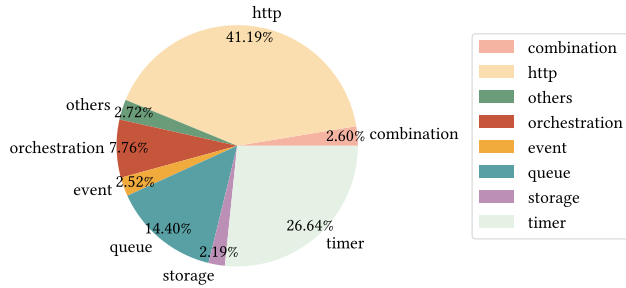


Fig. 5. The proportion of trigger types among functions.

Second, some triggers lead to regular invocation behaviors, offering possible clues for invocation prediction. We use the Kolmogorov-Smirnov test [37] to check if a single function’s invocations with more than 10 counts follow a given distribution ($pValue \geq 0.05$, *i.e.* not rejecting the null hypothesis). Results show that 68.12% timer-triggered functions are invoked periodically or quasi-periodically, excluding 6.65% do not have enough counts. Also, 45.02% HTTP-triggered invocations follow a Poisson arrival process, conforming to the expectation of human-generated activities, excluding 36.20% functions without enough samples. Though the Poisson process is memoryless and unsuitable for direct invocation prediction, it still indicates that meaningful invocation behaviors are associated with triggers.

2) *Application workflow*: An application contains several functions in a logic workflow [27], such as function chaining or callback calling, leading to recurring and interdependent invocation behaviors. Records also show that invocations from one application/user tend to behave consistently. We create a metric *co-occurrence rate (COR)* to measure such consistency.

For each function, we record the time when it and another function are invoked, divided by its invoked time, as its COR. The idea is borrowed from linguistics, where co-occurrence is interpreted as indicating semantic proximity. For each function with at least another function (*candidate function*) sharing the same application/user, we negatively sample 50 no-application/user-overlapped functions. The average COR of candidate functions is about 4.6x compared to negative samples, which are 0.2312 and 0.0504, respectively. Inspired by §III-B1, we also find that the invocation behavior convergence is more significant when candidates share the same trigger. Particularly, the average COR rate of the same-trigger-based candidate functions is 0.2710, whereas it is 0.1307 for different-trigger-based candidates.

3) *Temporal locality in invocations*: One more noteworthy behavior is temporal locality, the tendency of functions to experience short bursts in a short-term period. On the one hand, the triggering events can have a rapid succession due to real-world bursty requests (e.g., a high-season e-commerce site). On the other hand, many common application patterns contribute to temporal locality. In a fan-out application, a function activates multiple other functions concurrently, often resulting in rapid and successive calls to the secondary functions. Callback functions that handle asynchronous operations can also undergo multiple invocations in quick succession until the workload is completed. Such temporal locality also exists in the dataset. Figure 6 plots the minimum-maximum normalized invocations of five infrequently triggered functions, whose invocations are consecutive and concentrated in certain periods. For infrequently invoked functions presenting temporal locality, invocation prediction can be difficult, yet simply keeping a loaded function for some time can considerably reduce cold starts without wasting much memory.




Fig. 6. Invocations of these infrequently invoked five functions are centralized over several periods. We omit the no-invocation period (4000–13000).

Insight: The microservice and event-driven architecture of serverless systems results in the concentration of recurring and systematic invocation behaviors. By analyzing the invocation patterns and categorizing functions, we can accurately predict and pre-load function instances to reduce cold starts.

This insight inspires SPES, a rule-based, efficient, and elastic function scheduler. It leverages the application, user, and trigger information to design tailored and evolution-adaptive strategies to pre-load/evict categorized functions without requiring large amounts of training data. Overall, SPES provides a language-agnostic, developer-free function provision approach that can be adopted on existing FaaS systems seamlessly to optimize cold starts and memory usage.

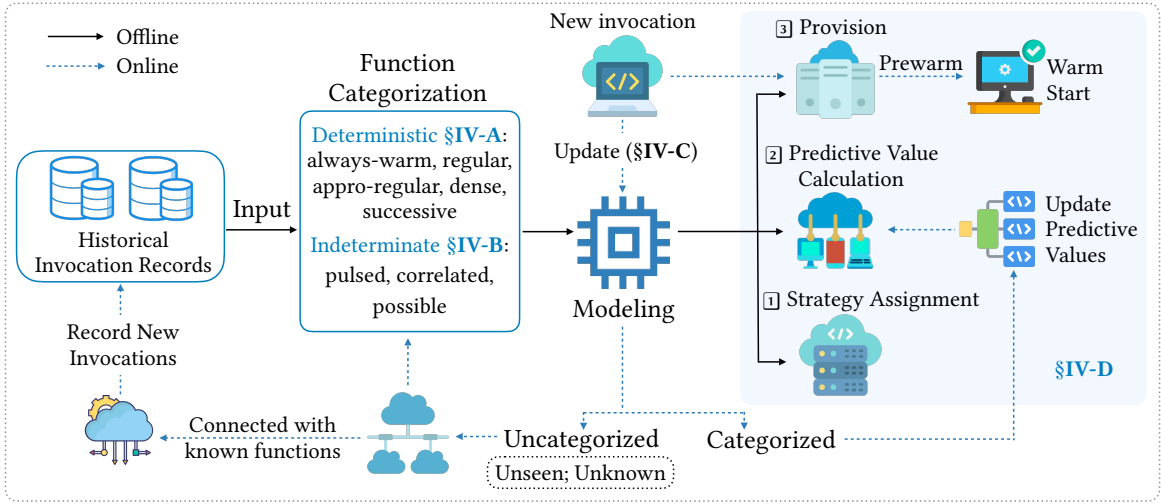


Fig. 7. Overview of SPES. In the offline phase, SPES models the historical invocation records for each function and categorizes functions according to their mined invocation patterns. During the online phase, if the new invocation comes from a categorized function, its function’s provision directly follows its provision rule, and uncategorized functions are tried to be connected with known ones. The invocation prediction and predictive values/indicators are then updated upon new invocations.

IV. METHODOLOGY

This section introduces the design of SPES. An ideal scheduler should decide to load a function exactly before its invocation and evict it from memory after the execution if no more invocations are imminent. The decision-making relies on fine-grained invocation prediction. To this end, we propose SPES, whose overview is presented in Figure 7. SPES consists of four parts: deterministic function categorization, indeterminate function assignment, adaptive strategy application, and function provision based on invocation prediction. Specifically, we first summarize typical invocation patterns, which result from pre-defined triggers [38] and the combinatorial methods of function calling (in common application scenarios [39]) and then formulate their corresponding definitions. If a function satisfies the definitions, it is categorized; Otherwise, the indeterminate functions are assigned to three supplementary types. We also design two extra adaptive strategies to handle the concept shifts as FaaS evolves. Finally, SPES provisions functions according to rule-based invocation prediction, where each type follows its own predicting rule.

Note that different triggers can exhibit the same timing feature in invocations. For example, the service bus trigger and the event gird trigger both can handle moving datagrams (though focusing on different message types). It is neither necessary nor desirable to correspond each trigger to a function type in the context of the cold start problem.

Let us start with three definitions:

- *Waiting time (WT)*: the length of successive idle time. Take an invocation sequence (28, 0, 12, 1, 0, 0, 0, 7) as an example, where each value denotes the invocation count at each sampling slot, derived to a sequence of WTs $\{WT\} = (1, 3)$, as no invocation exists at the 2nd slot or in slots 5–8. If we foreknow the next WT, for example, at the 4th slot, we know no invocations will arrive in slots 5–8, then we can easily make a perfect decision: evict the function now and re-load it at the end of slot 9. Hence, we can

transform invocation prediction to WT prediction. Different from a previously proposed inter-arrival time (IAT) [13], WT depicts the slot-grained interval between two successive invocation sequences, from the last’s end to the next’s start. WT is also different from idle time (IT). IT is the time without invocations or subsequent function executions, so IT exists after any single execution. WT, instead, only appears when a successive series of invocations ends, and the end of a single execution does not necessarily incur a WT.

- *Active time (AT)*: the length of successive time with invocations. The above-mentioned invocation sequence delivers an AT sequence $\{AT\} = (1, 2, 1)$ because the function is invoked in slots 1, 3–4, and 9.
- *Active number (AN)*: a descendable definition, the number of successive invocations during AT. Again, the above sequence delivers $\{AN\} = (28, 13, 7)$.

A. Deterministic Function Categorization

This section defines five invocation types based on typical invocation patterns and categorizes functions accordingly. Table I shows an overview of the five types. The guiding principle for definition is from easy to difficult, also serving as the categorization priority: if a function fits a former type, it will not fit any latter type. Our introduction follows this order.

1) *Always warm*: This type describes consistently active functions, such as long-running operations and hyperfrequent calls, which usually involve durable or timer-triggered functions (whose interval is small enough). For example, functions in continuous integration/deployment (CI/CD) pipelines are consistently invoked to automate code building, testing, and deployment. A function is defined as “always warm” if 1) it is invoked at every sampling time or 2) the sum of inter-invocation time is \leq one-thousandth the observing time. Such functions are undoubtedly always loaded.

2) *Regular*: Regular task processing, usually using timer-triggered functions, is a basic functionality (e.g., polling). Yet,

TABLE I
A BRIEF OVERVIEW OF THE WELL-DEFINED FUNCTION TYPES.

Type name	Characteristics	Definition	Predictive values
Always warm	Almost invoked all the time	Invoked at every time Inter-invocation time \leq observing time%	–
Regular	Almost invoked periodically	(Processed) $P_{95}(\{WT\}) - P_5(\{WT\}) \leq 1$ CV of $\{WT\} \leq 0.01$	Median of WTs
Appro-regular	Invoked quasi-periodically	(Processed) the count of $\text{Mode}_n(\{WT\}) \geq 0.9 \times$ sequence length	$\text{Mode}_n(\{WT\})$
Dense	Frequently invoked	$P_{90}(\{WT\}) \leq$ a small constant	$[\min, \max]$ of $\text{Mode}_k(\{WT\})$
Successive	Successively invoked	$\min(\{AT\}) \geq \gamma_1; \min(\{AN\}) \geq \gamma_2$ $\gamma_1 < \gamma_2$	–

their actual invocations may not be strictly periodic (with a constant WT) since there are undesired fluctuations: 1) The exact first/last WTs during an observing period are almost impossible to record. 2) Periodically generated events can be blocked or delayed by contingencies such as concurrency limits, network connectivity issues, etc. 3) Other events can invoke a mostly regularly invoked function occasionally.

Here comes the slacked definition: A “regular” function satisfies either 1) the difference between the 5th and 95th percentile of its WT sequence is ≤ 1 or 2) the coefficient of variation of WTs is close to zero (≤ 0.01 in practice). Otherwise, we remove the first and last WTs and re-check if the function follows the definition. If the answer is still no, we apply another slacking rule by merging adjacent small WTs. In particular, for each WT closely valued to the WT mode, its adjacent small WTs are gradually merged until reaching ① the sequence’s end or ② another WT close to the mode or ③ an already merged WT. In this way, a WT sequence valued by (1439, 1438, 1, 1439, 1438, 1) is processed to be (1439, 1439, 1439, 1439), and then it seems “regular”. With a processed WT sequence satisfying the above definition, the function belongs to the “regular” type. We partially eliminate accidental factors by processing WTs with slacking rules and modeling essential invocation behaviors. Finally, we record the median of WTs as the *predictive value* for invocation prediction, leveraged in Section IV-D.

3) *Approximatively regular*: This is the derived type from “regular”. A regular function may experience invocation variance and long-term disturbance. For example, a data-processing function is expected to receive updates from a data station every three minutes (IoT Hub trigger functions). Yet, limited by the data transmission capability, the function is actually invoked every 3–5 minutes. We define such functions as “approximatively (appro-)regular”. Particularly, a function is “appro-regular” if the occurring count of the first n frequently appearing values (modes) of WTs is \geq a large percent (90% practically) of the WT sequence’s length. n is a pre-defined integer. The predictive values for these “appro-regular” functions are the first n modes.

4) *Dense*: This characterizes irregular but frequent invocations with intermittent idle durations. “Dense” functions can involve several triggers: queues (or service bus using the queue

flavor) for frequent messages in asynchronous processing, influenced by variations in message arrival times; Cosmon DB triggers for frequent real-time data processing; HTTP triggers for processing frequent HTTP requests. Functions with a 90th percentile of the WTs \leq a small constant are defined as “dense.” The range of the first k modes derives predictive values, *i.e.* $[\min(\text{Mode}_k(\text{WTs})), \max(\text{Mode}_k(\text{WTs}))]$, where k is an empirical integer. Such a function should be unloaded only if its idle time is larger than the constant.

5) *Successive*: This describes inactive functions experiencing consecutive invocations in a short timeframe (temporal locality) before returning to an inactive state. It can involve HTTP trigger functions (e.g., bursty social media trends), storage trigger functions (e.g., uploading files to Blob Storage), etc. In practice, feedback loops engaged functions, caching mechanism implementations, and load balancing-involved applications can also exhibit such behavior. We define this as $\min(\{AT\}) \geq \gamma_1$ or $\{AN\} \geq \gamma_2$, where $\gamma_1 < \gamma_2$ and they are pre-defined lower bounds. Predicting the start of such an invocation wave is highly difficult. Nevertheless, we can tolerate the first invocation and keep the function alive until the wave ends. With limited cold starts, we save lots of memory.

B. Indeterminate Function Assignment

This section handles functions that do not satisfy the above five definitions using two strategies: ① “forgetting” previous records to re-check whether the not-too-long history invocations satisfy the five definitions, and ② assigning the remaining functions with three extra types.

1) *Forgetting*: As discussed in §III-A4, function invocations experience concept shifts. For example, a regular invoked function can be categorized as “always warm” if the interval is adjusted to be smaller than the recording unit. Therefore, the near-the-present records should be given priority. We slice out the sequential invocation observations in days, from the beginning to the half. Then, we gradually remove the older observations from the original d -day observations and check whether invocations in 2nd– d th days conform to existing definitions. If not, we check the invocations in 3rd– d th days, and so on, until the $\lfloor d/2 \rfloor$ th day.

2) *Assigning*: The remaining indeterminate functions are assigned to three types according to the validation results:

- D1. *Pulsed* describes functions showing less obvious temporal locality in invocations than “successive” ones. Similarly, we allow a cold start at the first invocation time and keep the function warm until its ideal time reaches a pre-defined threshold.
- D2. *Correlated* invocations are predicted based on other functions. Functions usually work in a logic workflow topology, interacting with other functions. This is present in several common application patterns, including function chaining (where a sequence of functions executes in a specific order), fan-out (where a function calls multiple functions in parallel), and fan-in (where a function waits for multiple functions to finish) [39]. We define such functions using T -lagged co-occurrence rate (T -COR), derived from COR in §III-B. T -COR is the COR between the target invocation sequence and an T -time-lagged invocation sequence of its candidate functions (those sharing an application/user). The lagged COR can better reflect how the invocation of candidate functions indicates near-future invocations of the target function. If the T -COR reaches a threshold (practically 0.5, $T \leq 10$), the candidate invocation is an important predictive indicator, and these two functions are linked. A function can have multiple linkages with different functions, including categorized and uncategorized ones.
- D3. *Possible* functions are infrequently invoked but have at least one mode of WTs (appears more than once). By taking the modes as predictive values, it is possible to predict invocations, though such prediction is deemed to be insufficiently satisfactory.

If any of the three strategies yields both the least cold starts and the minimum wasted memory during validation, the indeterminate function is directly assigned accordingly. Otherwise, the rate of rise is decided. Denote the cold starts of the above three strategies as (cs_1, cs_2, cs_3) and the corresponding wasted memory are wm_1, wm_2, wm_3 . Assume definition D1 delivers the minimum cold starts while strategy D2 delivers the minimum wasted memory. We compute the rise rates: $\Delta cs = (cs_2 - cs_1)/cs_1$ and $\Delta wm = (wm_1 - wm_2)/wm_2$. If $\Delta cs \times \alpha \leq \Delta wm$, then the function is assigned to D1, *i.e.* “plused”; otherwise, D2 prevails, where $\alpha \in (0, 1)$ is a scaling factor, and the smaller $\alpha \in (0, 1)$, the more importance is put on cold starts than wasted memory.

Most functions have been categorized so far, except those that have never been invoked in the validation. We simply leave them as “unknown” since they are likely very infrequently invoked and hardly have meaningful predictive indicators. Some unknown functions may be categorized during the online provision based on our adaptive strategies illustrated in the next section (§IV-C).

C. Adaptive Strategy Application

This section designs two adaptive strategies to handle concept shifts during the online provision: one adjusts the predictive values; the other is based on inter-function correlation. These strategies can further address some unknown

or unseen functions (those that never appear in the training data) with meaningful patterns during the online provision. The application of these strategies is also presented in Figure 7, where the invocation prediction and the prediction values are adaptively updated upon online invocations.

1) *Adjusting*: This adaptive strategy contains three steps:

- S1. Record the WTs during the online provision. If there are enough WTs, then initiate the adaptive updating.
- S2. Update the predictive values if the corresponding values computed from newly collected WTs change significantly with the mean of the old and new ones.
- S3. If the new WTs for an unknown or unseen function conform to previous definitions with enough samples, categorize the function as the corresponding type.

S2 requires more explanations. We record the predictive values for four types: “regular”, “appro-regular”, “dense”, and “possible”. Their predictive values are the median, the first n modes, the range of the first k modes, and the WT values occurring more than twice, respectively. Take “regular” as an example. Suppose the absolute difference between the median of offline WTs (*i.e.* the predictive value) and that of online WTs is larger than the standard of offline WTs. The new predictive value is updated by the mean of the old median and the new median. Other types adopt a similar adjusting strategy.

2) *Online correlation*: This strategy correlates unseen functions with known functions or appeared unseen functions, as the “correlated” strategy does. To speed up the computation and obtain informative indicators, we only consider candidate functions sharing the same trigger with the target function (*aka.* unseen function). Initially, if one of the candidates is invoked, we also pre-load the target function. Afterward, we gradually remove less correlated candidates. Again, we adopt COR to measure the degree of correlation. We count the pair-wise target-candidate COR at each slot and record the maximum. If the difference between a COR and the maximum is large enough, the corresponding candidate is kicked out from consideration unless its COR returns close to the maximum.

D. Next Invocation Prediction and Provision

This section predicts invocations and pre-loads functions based on the predictive values or indicators. “(Appro-)regular” functions have discrete predictive values and “dense” functions use continuous ones. In terms of “possible” functions, if the range of predictive values is larger than a threshold, the values are regarded as discrete; otherwise, we consider continuous integers inside the range of predictive values. The predicted invocation times are naturally the last invoked time added by each of the predictive values.

For the online provision, if one of the predicted invocation times falls in $[t - \theta_{prewarm}, t + \theta_{prewarm}]$ with a pre-defined parameter $\theta_{prewarm}$ at the time of t , the function will be pre-loaded. If a loaded function’s current WT is $\geq \theta_{givenup}$, the function will be evicted from the memory, where the parameter $\theta_{givenup}$ differs among function types. Algorithm 1 shows our provision optimization, which is rule-based with good

scalability and implements our defined strategies faithfully. Every time the algorithm returns the updated MemSet, based on which we load function instances.

Algorithm 1: The provision algorithm of SPES.

Input: FList: hash ids for all functions; FState: recording the necessary information about all functions, such as the current WT, history WTs, predictive values, the time of the last invocation, etc; Invo^(t): invocation numbers of each function at the time of t ; MemSet: hash ids of loaded functions; UCorr: the correlations of unseen functions.

Output: The updated MemSet, FState and UCorr

```

1 Function Provision (FList, FState, Invo(t), MemSet,
  UCorr):
2   for  $f \in FList$  do
3     if  $Invo^{(t)}[f] > 0$  then
4       FState[f].last_invoked  $\leftarrow t$ ;
5       FState[f].update_WTs(FState[f].current_WT);
6       FState[f].current_WT  $\leftarrow 0$ ;
7       // Adaptively adjusting predictive values.
8       FState[f].update_predictive_values(t);
9       if  $f \notin MemSet$  then
10        FState[f].cold_start_record(t);
11        MemSet.add(f);
12      end
13    else
14      FState[f].current_WT += 1;
15      pre_load_flag  $\leftarrow$ 
16        FState[f].pre_load(FState[f]. $\theta_{prewarm}$ );
17      if pre_load_flag is False AND
18        FState[f].current_WT  $\geq$  FState[f]. $\theta_{givenup}$ 
19      then
20        MemSet.remove(f);
21      else if pre_load_flag is True then
22        MemSet.add(f); // Pre-loading.
23    end
24    // Adaptively processing unseen functions.
25    UCorr.update();
26  end
27  return MemSet, FState, UCorr
28 End

```

V. EVALUATION

We evaluate SPES by answering four research questions:

- **RQ1:** How effectively does SPES decrease cold starts?
- **RQ2:** How much memory waste and computation overhead does SPES incur?
- **RQ3:** How does SPES trade off memory waste with latency reduction?
- **RQ4:** How do complementary designs influence SPES?

A. Experiment Settings

We evaluate SPES on the most widely used industrial dataset [40] released by Microsoft Azure Function [13] with real-world invocation traces. The dataset contains the invocation counts per minute for 14 days. The first 12 days are used for pattern modeling (training), and the last two days are used for the simulation, conducted on a workstation with an 8-core Intel i5-3470S CPU and 16 GB memory.

The records involve 15,097 users, 24,964 applications, and 83,137 functions. 71,616 functions appear in the training data, 39,388 functions appear in the simulation data, and 743 never appeared during training.

We adopt the simulation principles following [13]. First, assume all executions finish within one minute since 1) most (96%) functions have very short execution time (less than 60 seconds), and 2) we can thereby calculate the worst-case wasted resource time. Second, we assume that cold-start latency for each function is the same, so we only need to care about the number of cold starts and the wasted memory time.

1) *Baselines:* We compare SPES with five state-of-the-art baselines applied to the application layer. Approaches involving system renovations are out of this paper’s scope and will be discussed in Section VII. Our baselines include FaaSCache [15], Defuse [22], Hybrid [13], and a fixed keep-alive policy. We reproduce these methods as they do not provide open-source code by strictly following the original papers and referring to a reproduction attempt [41]. Note that the original Hybrid method works at the application unit (Hybrid-Application, HA), so we derive a Hybrid-Function method by employing its design on the function granularity (Hybrid-Function, HF), following [22]. All the parameters are set according to their original papers. The fixed keep-alive policy adopts a length of 10 mins. FaaSCache requires a pre-defined memory limit, so we adopt the maximum memory size of SPES during the whole simulation.

2) *Metrics and parameters:* To quality the cold-start optimization, we measure the function-wise (application-wise for HA) *cold-start rate (CSR)*, i.e. the number of cold starts divided by the number of invocations. We apply wasted memory time (*WMT*, §II-B) to gauge the idle resource waste. Naturally, the lower the CSR or WMT, the fewer cold starts or wasted resources, the better. We also monitor the *effective memory consumption ratio (EMCR)*, which measures the fraction of invoked function instances relative to the total loaded instances on each host machine, serving to assess resource efficiency. A higher EMCR signifies wiser memory allocation, as it indicates a greater proportion of memory is used by active instances rather than remaining idle. As for the parameters, we set $\theta_{prewarm}$ as two. The $\theta_{givenup}$ for “dense” and “plused” is five, whereas one for the other types. Section V-D further discusses the impact of these pre-defined parameters.

B. RQ1: Effectiveness in Cold-Start Reduction

Figure 8 displays the cumulative distribution (CDF) of CSR under the provision decisions of SPES and baselines. With a fixed y-axis value, the line representing SPES is positioned to the left of other baseline lines, which indicates that SPES consistently leads to fewer cold starts across a variety of functions, each corresponding to different invocation frequencies. In particular, SPES reduces the 75th percentile cold-start rate (Q3-CSR, for simplicity) from 0.215 to 0.108 compared to Defuse, the best-performing baseline, achieving 49.77% improvement, and reduces 75-CSR by 64.06%–89.20% compared to other baselines. We care more about Q3-

CSR because infrequently invoked functions benefit most from optimization [13]. Moreover, 57.99% functions experience no cold starts with SPES, indicating that SPES can allow most functions to be warmly invoked. In contrast, 25.61%–52.59% functions completely experience warm startup using baselines where FaaSCache is the best one. On the other hand, regarding infrequently invoked functions, SPES reduces the 90th percentile CSR by 19.87% compared to the best-performing baseline, HA. Such improvements imply that SPES exhibits significant optimization performance for both frequently and infrequently invoked functions, while no baseline can achieve second-best performance with different percentile CSRs.

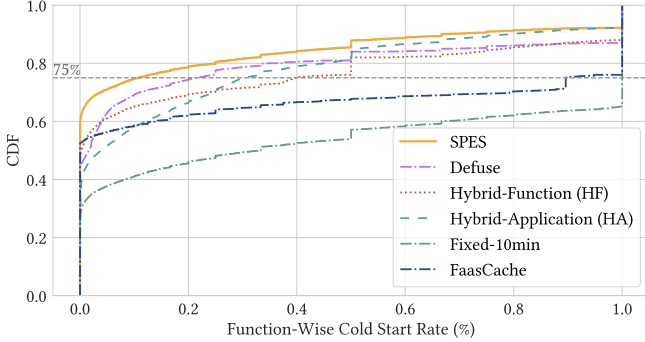


Fig. 8. Cold start behaviors of SPES and its competitors.

Furthermore, SPES significantly reduces “always-cold” functions, which always experience cold starts upon invocations (with CSR equal to 1.0). Figure 9(b) further presents the always-cold percentage, and that of SPES is just less than 8%. Among the baselines, HA has the fewest always-cold functions, which is closest to that of SPES, whereas Defuse and HF, both function-grained methods, increase always-cold functions dramatically. This can be attributed to infrequently invoked functions. About 3.82% functions are invoked less than twice during training, and 6.14% are only once during the simulation, so always-cold functions seem inevitable with limited records. HA mitigates this issue by grouping and loading functions together, yet resulting in more memory consumption and waste. SPES instead connects unseen and unpredictable functions with known functions using trigger, application, and user information and gradually updates the associated functions based on actual invocations. Hence, without much extra memory, SPES performs close to HA.

In addition to cold start reduction, Figure 9(a) presents the memory usage normalized by the averaged one of SPES. SPES’s memory usage is only 8.08% more than the most resource-efficient method, the fixed keep-alive policy, and saves 36.07%–55.55% memory compared to the other baselines on average. It only consumes about half of the memory of Defuse, the best baseline for cold start reduction.

Figure 10 shows the averaged CSR among different types. As the simulation period is not long enough, we only categorize unknown functions into the type of “possible”, denoted by “newly-possible”. We can see that “unknown” functions contribute the most to cold starts, and “pulsed” functions also incur high CSRs. This is attributed to insufficient historical




Fig. 9. SPES is memory-efficient compared to most baselines with relatively few always-cold functions.

invocation records. Actually, SPES intentionally convives a cold start when an “unknown” or “pulsed” invocation arrives after a long idle time. Though we can leverage less predictive information, such as the averaged WT, to predict the next invocation or even keep such functions always warm, this can lead to considerable unbillable and wasted memory, undesired for FaaS providers. This outcome is deliberating on the trade-off between performance and resource allocation. Such an issue can be mitigated with more invocation histories revealing predictive indicators for better provision.





Fig. 10. Averaged cold start rate of each type.

Meanwhile, though “correlated” and “possible/newly-possible” functions also lack predictive information with infrequent invocations, their invocations are more likely to be predicted. Thus, we connect some of them with deterministic functions or adaptively extract meaningful behavior indicators from incoming invocations. In this way, their cold starts are suppressed effectively. We will further discuss the usefulness of such strategy designs in Section V-E.

C. RQ2: Wasted Memory Time and Overhead

1) *Wasted Memory Time*: Figure 11(a) presents that SPES significantly decreases the wasted memory time by 10.89%–63.50% compared to all baselines. Particularly, compared to Defuse, the most effective baseline on cold-start reduction, we reduce 57.06% of the WMT. Figure 11(b) also demonstrates that SPES efficiently uses memory resources, whose EMCR is 46.32%, 5.20%–120.89% higher than compared approaches. The success of SPES can be attributed to careful pattern modeling and differentiated strategies.

To further demonstrate how each type contributes to WMT, we derive a metric called the *ratio of WMT*, which is the WMT divided by the number of invoked times for each serverless function. If a type contains more functions or function invocations, the bespoke scheduling strategy tends



(a) Normalized wasted memory time (WMT, lower is better) (b) Effective memory consumption ratio (EMCR, higher is better)

Fig. 11. SPES significantly reduces memory waste and achieves more effective memory consumption.

to make more proactive provisions, likely resulting in more WMT. Thus, the ratio of WMT can better portray the accuracy of invocation prediction and the usefulness of pre-loading. Figure 12 shows the distribution of the ratio of WMT among different function types, from which we can see the “possible” functions have the highest probability of generating WMT. The predominant reason for this is the infrequent invocation of possible functions, which leads to a scarcity of patterns in the historical database and complicates establishing associations with other functions. Despite these challenges, we persist in our efforts to forecast and pre-warm the invocation of these functions, aiming to reduce cold starts as much as possible. Different from “pulsed” or “unknown” functions that are allowed to generate cold starts, we encourage aggressive prediction attempts for “possible” functions since the latter have at least a duplicated WT, enabling potential predictive value obtaining. However, this strategy unavoidably results in augmented resource wastage.




Fig. 12. The ratio of WMT of each function type.

2) *SPES Overhead*: Herein, we measure the additional latency induced by our implementation. The time complexity within each time window of our approach and the fixed keep-alive policy is $\mathcal{O}(n)$, where n represents the number of functions, as it only requires accessing the corresponding function’s category and its predictive value(s). For each invocation, the time complexity of pre-load/unload operations is merely $\mathcal{O}(1)$. Within each time window, the time complexity for the fixed keep-alive method is $\mathcal{O}(m)$, where m is the number of loaded functions. In contrast, FaaSCache has a time complexity of $\mathcal{O}(\log m)$, which involves identifying the container with the lowest priority. Other methods, such as HA, HF, and Defuse, have higher complexities, mainly due to the computational bottleneck in updating histograms. The simulation results align with these expectations. Compared to the fastest fixed keep-

alive baseline, which has an average overhead of 0.024 sec per minute, our method adds 0.44 sec’s overhead per minute, mainly due to our more complex strategies for each provision action. Contrasted with the second fastest, FaaSCache, our overhead is reduced by 6.8%. In summary, our overhead is inconsequential compared to the typical latency found in most existing serverless platforms.

D. RQ3: Trading-off resources and latency

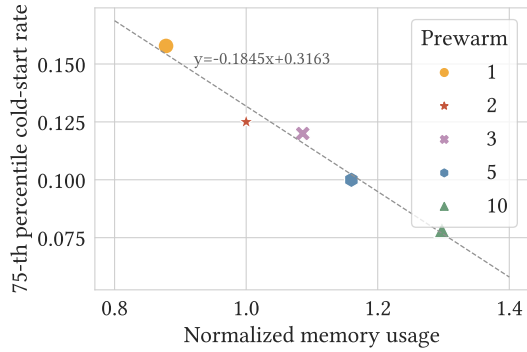
We control the trade-off through two parameters: $\theta_{prewarm}$ and $\theta_{givenup}$. As introduced in Section IV-D, the former decides how long to pre-load a function with a predicted nearby invocation, and the latter decides how long to keep an idle function warm. Intuitively, the larger these two parameters, the more likely a function is to pre-load or keep loaded with more memory usage and potential wasted memory, and the less cold starts. In RQ1 (§V-B), we set $\theta_{prewarm}$ as two, and the $\theta_{givenup}$ for “dense” and “plused” is five, whereas one for the other types. Our original simulation setting is denoted by the red star (\star) in the following figures.

Figure 13(a) shows the trade-off under different $\theta_{prewarm}$, where a point (x, y) represents using x -unit memory and obtaining the 75-CSR of y , under a certain $\theta_{prewarm}$. The memory is normalized to that under the original setting. The normalized memory usage and the 75-CSR are nearly linearly correlated. We can conveniently choose a proper setting by controlling $\theta_{prewarm}$. Since the red star is below and the most distant from the fitting line, $\theta_{prewarm} = 2$ is the optimal value.


As $\theta_{givenup}$ should be integers and different from each type, we simply multiply the original $\theta_{givenup}$ setting by 2, 3, 4, and 5, respectively. The results are shown in Figure 13(b), where the linear relationship still approximately holds, but larger $\theta_{givenup}$ s have less impact on cold start reduction. This indicates that keeping invoked functions too long is sub-optimal and idle functions should be evicted promptly.

E. RQ4: Impact of Strategy Designs

1) *Impact of inter-function correlations*: This section gauges our design in processing ill-informed functions by developing inter-function correlations. As introduced in §IV-B, we propose a simple yet effective T -lagged co-occurrence rate metric to connect ill-informed functions with categorized ones. Those with closely related known functions are “correlated”. This strategy is applied during both training and simulation. Figure 14 presents the impact of this strategy. *w/o Corr* means re-categorizing “correlated” functions into “pulsed”, “possible” or “unknown” during training, but we still deal with unseen functions during the simulation. *w/o Online-Corr* denotes removing the simulation-applied strategy, *i.e.* regarding unseen functions as “unknown” but still retaining the “correlated” functions obtained during training. It is shown that the latter strategy slightly reduces the Q3-CSR, whereas the former makes a significant contribution. We attribute the results to the influenced function number. 4.71% of the functions belong to “correlated” whereas only 1.89% are unseen.



(a) Parameter $\theta_{prewarm}$



(b) Parameter θ_{given}

Fig. 13. Under different $\theta_{prewarm}$, the memory usage and Q3-CSR exhibit an approximately linear relationship. Similarly, under different θ_{given} , the linear relationship still approximately holds, yet dramatically increasing the memory does not guarantee cold start mitigation.

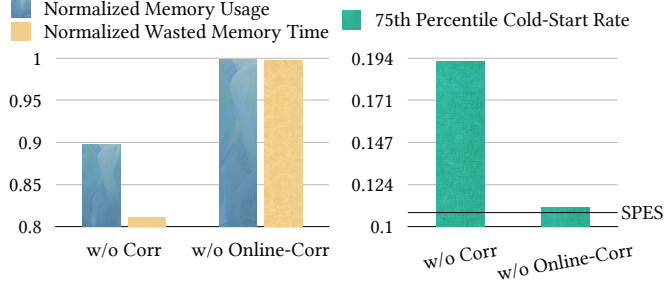


Fig. 14. The correlation strategy contributes to the cold-start and memory-waste reduction.

2) *Impact of designs regarding concept shifts*: This section studies the impact of two designs regarding concept shifts: 1) *forgetting*: forcing unknown functions to a defined type by ignoring older invocations while focusing more on recent data (§IV-B); 2) *adjusting*: adaptively adjusting the predictive values during the simulation (§IV-C). Figure 15 depicts the effectiveness after omitting these two adaptive designs, respectively. Removing the second shifting strategy has a slighter impact. Similarly, this is attributed to the fact that the forgetting strategy involves more function categorization efforts. We do not pre-load unknown functions, so the forgetting strategy categorizing 340 unknown functions has a larger impact. In contrast, the adjusting strategy only categorizes 174 unknown functions into the “newly-possible” type and updates the predictive values of 499 “(appro-)regular/dense” functions, resulting in less impact. Nevertheless, both designs contribute

to the effectiveness of SPES. These designs will play a greater value with more data and a longer simulation period (usually indicating larger shifts).

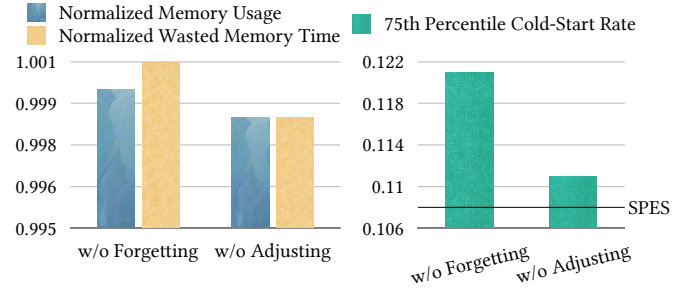


Fig. 15. SPES’ adaptivity benefits cold-start reduction.

VI. DISCUSSION

A. Limitations

1) *Idealized Simulation Principles*: Our simulation adopts two key assumptions to simplify the process and mitigate dataset limitations. Firstly, we assume each function execution completes within one minute. For functions with longer execution, adjusting the counting unit as necessary to accommodate the vast majority (e.g., 95%) of functions within this timeframe can ensure the assumption’s relevance. For functions exceeding this one timeframe without new invocations—e.g., for a two-minute execution without activity in the second minute—we maintain memory allocation and increment the WMT despite no memory being actually wasted. Such deviation can particularly underestimate resource efficiency, especially for methods with frequent evictions, including ours and HF. Thus, we posit that such marginal deviations do not affect the overall integrity of our experimental findings. Secondly, we assume all functions share the same cold-start latency. Our focus on function-specific metrics means that variations in cold-start latency across functions do not significantly affect our primary research insights, though a more detailed dataset could enhance our analysis of the aggregate cold-start latency and offer a deeper comparison of methodologies.

2) *Over-simplicity of memory consumption*: We simply assume that one node can accommodate all loaded function instances with infinite resources following [22], [13], so we do not consider problems like worker communication, sandbox optimization, and load balancing. Otherwise, we may need to renovate the underlying system intrusively. In addition, we regard all function instances as equally consuming the memory limited by the data that do not contain function-wise memory consumption. They actually require different memory allocations, yet the assumption is still reasonable as these functions are usually responsible for simple tasks with confined memory size [22], [42]. We leave these system-layer challenges to future work beyond this paper’s scope.

3) *Ignorance of the function priority*: We simply regard all functions with the same level of importance or urgency, but a real-world platform may prioritize functions to ensure that time-sensitive or mission-critical workloads are executed promptly, even during periods of high demand or resource

constraints to meet QoS requirements [43]. We plan to improve SPES by adding a hierarchical module to dictate how to prioritize, schedule or queue functions in a QoS-aware manner.

B. Threats to Validity

1) *Internal*: The evolution of serverless functions is very fast. With only 14-day data, we can observe distinct concept shifts in function invocations. We currently do not conduct proactive re-categorization or rule generation. The performance of SPES may degrade in the very long term. To mitigate this threat, we design adaptive strategies to handle relatively small shifts, which can easily be extended to long-term shifts.

2) *External*: The effectiveness of SPES on other datasets is yet unknown. To alleviate this concern, we choose Azure Function, the most widely used dataset from real-world industry, whose representativeness has been confirmed by many studies [13], [22], [44], [45], [46]. Moreover, it is the only publicly available industrial dataset. We will strengthen cooperation with the industry to evaluate SPES on more datasets and prompt the implementation of SPES on a real-world FaaS system. We also release our code to facilitate reproductivity and more experiments in the community.

VII. RELATED WORK

Numerous efforts [35], [47], [48], [13], [22], [45], [19], [15], [49] have focused on minimizing cold starts and memory overhead for FaaS platforms. Related studies [17], [21] classify current methods into two categories: speeding up function warming (involving renovations in the system layer) and reducing occurrences of runtime cold starts (at the application layer). Methods of these two categories are orthogonal, so they can be combined for further cold start optimization.

This first category presents new systems to optimize infrastructure and resource management, employing techniques like load balancing among worker nodes [35], [47], snapshotting [50], [14], sandbox (e.g., container) scheduling and management [23], [48], [19], [51], [52], [53], [54]. One of the most popular research directions is sandbox sharing, i.e., sandboxes of the previous execution are reused for new invocations since reusing the idle sandbox incurs less delay than allocating a new one [51]. A recent study, Pagurus [52], proposed a container management scheme that allows one function's idle warm container to be forked by another function to alleviate cold starts, showing promising results. However, implementing these methods requires significant engineering efforts, system expertise, and ongoing maintenance due to the underlying platform and sandbox modifications, which pose challenges in diverse infrastructures and platforms.

The second category mitigates cold-start occurrences [13], [22], [15], [45], [23], [49] through a non-intrusive way. A primary step of function scheduling is to decide when and whether to pre-load/evict a function instance to reduce cold starts. [13] utilizes a small histogram to monitor function inter-invocation times, benefiting workloads with clear invocation patterns by optimizing keep-alive and pre-warming. However, it is fully data-driven and ignores the underlying invocation

patterns, leaving much room for domain-knowledge-involved invocation prediction. Defuse [22] employs dependency mining from function invocation histories to optimize the keep-alive time and pre-warming. Yet, it relies on the statistical histogram and turns to a fixed keep-alive policy for more than 32% of the functions, delivering inadequate perfection of effectiveness. LCS [49] selected the least recently warm container to reduce cold starts by keeping the containers alive for a longer period. FaaS-Cache [15] innovatively establishes an equivalence between keeping functions alive and keeping objects in a cache, thereby implementing function provision based on Greedy-Dual-Size Frequency object caching. However, its fundamental idea is to use up the given resources as much as possible until they are insufficient, then evict containers to minimize cold starts. Thus, its way of optimizing resources is rigid without any predictions, failing to smoothly optimize cold starts and resource usage simultaneously. This leads to a more significant waste of memory resources.

There is improvement space for these function provisioning methods as they lack an understanding of invocation patterns and fine-grained invocation prediction. Instead, SPES adopts horses-for-courses rule-based strategies to enhance existing approaches via accurate invocation prediction. SPES also proposes adaptive designs and builds up inter-function correlations, delivering better memory-economic cold-start reduction. Other efforts target a downstream task to determine a suitable node for scheduling an individual function request. FaaS-Rank [45] and ENSURE [23] attempt to pack load on the adequate number of invokers, allowing the additional unneeded invokers to idle, so as to reduce the completion time and cold starts of functions. Our method can be combined with these methods for finer-grained scheduling.

VIII. CONCLUSION

This paper proposes the first differentiated runtime serverless function provision method for cold start optimization. Our insight is that typical invocation patterns resulting from the underlying serverless architecture and triggers with commonness can lead to predictive invocation behaviors. Hence, we propose SPES, which provisions functions based on categorization and invocation prediction using tailored strategies. Our rule-based solution effectively mitigates the cold start problem by addressing the challenges of efficiency, scalability, imbalance, and evolution. Experiments demonstrate that SPES achieves significant improvement on both sides of the trade-off: economizing expensive resources and reducing cold starts to mitigate latency. Lastly, we make our code publicly available to facilitate future research and deployment of FaaS.

ACKNOWLEDGEMENT

This paper was supported by the Guangdong Key Research Program (No. 2020B010165002), the Research Grants Council of the Hong Kong Special Administrative Region, China (No. CUHK 14206921 of the General Research Fund), Shenzhen Stability Science Program, and the National Natural Science Foundation of China (No. 62202511).

REFERENCES

- [1] K. Namee, R. Phoarun, G. M. Albadrani, J. Polpinij, S. Tanessakulwattana, and P. Sphanphong, "A form and API data management platform for progressive web application and serverless application architecture," in *CIIS 2019: The 2nd International Conference on Computational Intelligence and Intelligent Systems, Bangkok, Thailand, November, 2019*. ACM, 2019, pp. 144–149. [Online]. Available: <https://doi.org/10.1145/3372422.3372452>
- [2] R. Kesavan, D. Gay, D. Thevessen, J. Shah, and C. Mohan, "Firestore: The nosql serverless database for the application developer," in *39th IEEE International Conference on Data Engineering, ICDE 2023, Anaheim, CA, USA, April 3-7, 2023*. IEEE, 2023, pp. 3376–3388. [Online]. Available: <https://doi.org/10.1109/ICDE55515.2023.00259>
- [3] Y. Yang, L. Zhao, Y. Li, H. Zhang, J. Li, M. Zhao, X. Chen, and K. Li, "Inflex: a native serverless system for low-latency, high-throughput inference," in *ASPLOS '22: 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Lausanne, Switzerland, 28 February 2022 - 4 March 2022*. ACM, 2022, pp. 768–781. [Online]. Available: <https://doi.org/10.1145/3503222.3507709>
- [4] J. Carreira, P. Fonseca, A. Tumanov, A. Zhang, and R. H. Katz, "Cirrus: a serverless framework for end-to-end ML workflows," in *Proceedings of the ACM Symposium on Cloud Computing, SoCC 2019, Santa Cruz, CA, USA, November 20-23, 2019*. ACM, 2019, pp. 13–24. [Online]. Available: <https://doi.org/10.1145/3357223.3362711>
- [5] H. Yu, H. Wang, J. Li, X. Yuan, and S. Park, "Accelerating serverless computing by harvesting idle resources," in *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*. ACM, 2022, pp. 1741–1751. [Online]. Available: <https://doi.org/10.1145/3485447.3511979>
- [6] S. Gupta, S. Rahnema, E. Linsenmayer, F. Nawab, and M. Sadoghi, "Reliable transactions in serverless-edge architecture," in *39th IEEE International Conference on Data Engineering, ICDE 2023, Anaheim, CA, USA, April 3-7, 2023*. IEEE, 2023, pp. 301–314. [Online]. Available: <https://doi.org/10.1109/ICDE55515.2023.00030>
- [7] P. Datta, P. Kumar, T. Morris, M. Grace, A. Rahmati, and A. Bates, "Valve: Securing function workflows on serverless computing platforms," in *WWW '20: The Web Conference 2020, Taipei, Taiwan, April 20-24, 2020*. ACM / IW3C2, 2020, pp. 939–950. [Online]. Available: <https://doi.org/10.1145/3366423.3380173>
- [8] L. Ao, G. Porter, and G. M. Voelker, "Faasnap: Faas made fast using snapshot-based vms," in *EuroSys '22: Seventeenth European Conference on Computer Systems, Rennes, France, April 5 - 8, 2022*. ACM, 2022, pp. 730–746. [Online]. Available: <https://doi.org/10.1145/3492321.3524270>
- [9] A. Jangda, D. Pinckney, Y. Brun, and A. Guha, "Formal foundations of serverless computing," *Proc. ACM Program. Lang.*, vol. 3, no. OOPSLA, pp. 149:1–149:26, 2019. [Online]. Available: <https://doi.org/10.1145/3360575>
- [10] A. Agache, M. Brooker, A. Iordache, A. Liguori, R. Neugebauer, P. Piwonka, and D. Popa, "Firecracker: Lightweight virtualization for serverless applications," in *17th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2020, Santa Clara, CA, USA, February 25-27, 2020*. USENIX Association, 2020, pp. 419–434. [Online]. Available: <https://www.usenix.org/conference/nsdi20/presentation/agache>
- [11] R. Haas, R. Niedermayr, T. Roehm, and S. Apel, "Is static analysis able to identify unnecessary source code?" *ACM Trans. Softw. Eng. Methodol.*, vol. 29, no. 1, pp. 6:1–6:23, 2020. [Online]. Available: <https://doi.org/10.1145/3368267>
- [12] A. Eivy, "Be wary of the economics of "serverless" cloud computing," *IEEE Cloud Comput.*, vol. 4, no. 2, pp. 6–12, 2017. [Online]. Available: <https://doi.org/10.1109/MCC.2017.32>
- [13] M. Shahrads, R. Fonseca, I. Goiri, G. Chaudhry, P. Batum, J. Cooke, E. Laureano, C. Tresness, M. Russinovich, and R. Bianchini, "Serverless in the wild: Characterizing and optimizing the serverless workload at a large cloud provider," in *2020 USENIX Annual Technical Conference, USENIX ATC 2020, July 15-17, 2020*. USENIX Association, 2020, pp. 205–218. [Online]. Available: <https://www.usenix.org/conference/atc20/presentation/shahrads>
- [14] D. Ustiugov, P. Petrov, M. Kogias, E. Bugnion, and B. Grot, "Benchmarking, analysis, and optimization of serverless function snapshots," in *ASPLOS '21: 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Virtual Event, USA, April 19-23, 2021*. ACM, 2021, pp. 559–572. [Online]. Available: <https://doi.org/10.1145/3445814.3446714>
- [15] A. Fuerst and P. Sharma, "Faas-cache: keeping serverless computing alive with greedy-dual caching," in *ASPLOS '21: 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Virtual Event, USA, April 19-23, 2021*. ACM, 2021, pp. 386–400. [Online]. Available: <https://doi.org/10.1145/3445814.3446757>
- [16] C. Chen, L. Nagel, L. Cui, and F. P. Tso, "S-cache: Function caching for serverless edge computing," in *Proceedings of the 6th International Workshop on Edge Systems, Analytics and Networking, EdgeSys 2023, Rome, Italy, 8 May 2023*. ACM, 2023, pp. 1–6. [Online]. Available: <https://doi.org/10.1145/3578354.3592865>
- [17] P. Vahidinia, B. J. Farahani, and F. S. Aliee, "Cold start in serverless computing: Current trends and mitigation strategies," in *2020 International Conference on Omni-layer Intelligent Systems, COINS 2020, Barcelona, Spain, Aug 31 - Sept 2, 2020*. IEEE, 2020, pp. 1–7. [Online]. Available: <https://doi.org/10.1109/COINS49042.2020.9191377>
- [18] E. Oakes, L. Yang, D. Zhou, K. Houck, T. Harter, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau, "SOCK: rapid task provisioning with serverless-optimized containers," in *2018 USENIX Annual Technical Conference, USENIX ATC 2018, Boston, MA, USA, July 11-13, 2018*. USENIX Association, 2018, pp. 57–70. [Online]. Available: <https://www.usenix.org/conference/atc18/presentation/oakes>
- [19] D. Du, T. Yu, Y. Xia, B. Zang, G. Yan, C. Qin, Q. Wu, and H. Chen, "Catalyzer: Sub-millisecond startup for serverless computing with initialization-less booting," in *ASPLOS '20: Architectural Support for Programming Languages and Operating Systems, Lausanne, Switzerland, March 16-20, 2020*. ACM, 2020, pp. 467–481. [Online]. Available: <https://doi.org/10.1145/3373376.3378512>
- [20] Z. Li, Y. Liu, L. Guo, Q. Chen, J. Cheng, W. Zheng, and M. Guo, "Faasflow: enable efficient workflow execution for function-as-a-service," in *ASPLOS '22: 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Lausanne, Switzerland, 28 February 2022 - 4 March 2022*. ACM, 2022, pp. 782–796. [Online]. Available: <https://doi.org/10.1145/3503222.3507717>
- [21] X. Liu, J. Wen, Z. Chen, D. Li, J. Chen, Y. Liu, H. Wang, and X. Jin, "Faaslight: General application-level cold-start latency optimization for function-as-a-service in serverless computing," *ACM Trans. Softw. Eng. Methodol.*, vol. 32, no. 5, jul 2023. [Online]. Available: <https://doi.org/10.1145/3585007>
- [22] J. Shen, T. Yang, Y. Su, Y. Zhou, and M. R. Lyu, "Defuse: A dependency-guided function scheduler to mitigate cold starts on faas platforms," in *41st IEEE International Conference on Distributed Computing Systems, ICDCS 2021, Washington DC, USA, July 7-10, 2021*. IEEE, 2021, pp. 194–204. [Online]. Available: <https://doi.org/10.1109/ICDCS51616.2021.00027>
- [23] A. Suresh, G. Somashekar, A. Varadarajan, V. R. Kakarla, H. Upadhyay, and A. Gandhi, "ENSURE: efficient scheduling and autonomous resource management in serverless environments," in *IEEE International Conference on Autonomic Computing and Self-Organizing Systems, ACSOS 2020, Washington, DC, USA, August 17-21, 2020*. IEEE, 2020, pp. 1–10. [Online]. Available: <https://doi.org/10.1109/ACSOS49614.2020.00020>
- [24] J. Sampé, G. Vernik, M. S. Artigas, and P. G. López, "Serverless data analytics in the IBM cloud," in *Proceedings of the 19th International Middleware Conference, Middleware Industrial Track 2018, Rennes, France, December 10-14, 2018*. ACM, 2018, pp. 1–8. [Online]. Available: <https://doi.org/10.1145/3284028.3284029>
- [25] S. Burckhardt, B. Chandramouli, C. Gillum, D. Justo, K. Kallas, C. McMahon, C. Meiklejohn, and X. Zhu, "Netherite: Efficient execution of serverless workflows," *Proc. VLDB Endow.*, vol. 15, no. 8, pp. 1591–1604, 2022. [Online]. Available: <https://www.vldb.org/pvldb/vol15/p1591-burckhardt.pdf>
- [26] Y. Wu, T. T. A. Dinh, G. Hu, M. Zhang, Y. M. Chee, and B. C. Ooi, "Serverless data science - are we there yet? A case study of model serving," in *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*. ACM, 2022, pp. 1866–1875. [Online]. Available: <https://doi.org/10.1145/3514221.3517905>
- [27] V. M. Bhasi, J. R. Gunasekaran, P. Thinakaran, C. S. Mishra, M. T. Kandemir, and C. R. Das, "Kraken: Adaptive container provisioning

- for deploying dynamic dags in serverless platforms,” in *SoCC '21: ACM Symposium on Cloud Computing, Seattle, WA, USA, November 1 - 4, 2021*. ACM, 2021, pp. 153–167. [Online]. Available: <https://doi.org/10.1145/3472883.3486992>
- [28] J. Jiang, S. Gan, Y. Liu, F. Wang, G. Alonso, A. Klimovic, A. Singla, W. Wu, and C. Zhang, “Towards demystifying serverless machine learning training,” in *SIGMOD '21: International Conference on Management of Data, Virtual Event, China, June 20-25, 2021*. ACM, 2021, pp. 857–871. [Online]. Available: <https://doi.org/10.1145/3448016.3459240>
- [29] Z. Zhou, Y. Zhang, and C. Delimitrou, “AQUATOPE: qos-and-uncertainty-aware resource management for multi-stage serverless workflows,” in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1, ASPLOS 2023, Vancouver, BC, Canada, March 25-29, 2023*. ACM, 2023, pp. 1–14. [Online]. Available: <https://doi.org/10.1145/3567955.3567960>
- [30] A. Singhvi, A. Balasubramanian, K. Houck, M. D. Shaikh, S. Venkataraman, and A. Akella, “Atoll: A scalable low-latency serverless platform,” in *SoCC '21: ACM Symposium on Cloud Computing, Seattle, WA, USA, November 1 - 4, 2021*. ACM, 2021, pp. 138–152. [Online]. Available: <https://doi.org/10.1145/3472883.3486981>
- [31] L. Wang, M. Li, Y. Zhang, T. Ristenpart, and M. M. Swift, “Peeking behind the curtains of serverless platforms,” in *2018 USENIX Annual Technical Conference, USENIX ATC 2018, Boston, MA, USA, July 11-13, 2018*. USENIX Association, 2018, pp. 133–146. [Online]. Available: <https://www.usenix.org/conference/atc18/presentation/wang-liang>
- [32] C. Q. Adamsen, A. Møller, S. Alimadadi, and F. Tip, “Practical AJAX race detection for javascript web applications,” in *Proceedings of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2018, Lake Buena Vista, FL, USA, November 04-09, 2018*. ACM, 2018, pp. 38–48. [Online]. Available: <https://doi.org/10.1145/3236024.3236038>
- [33] A. Wang, S. Chang, H. Tian, H. Wang, H. Yang, H. Li, R. Du, and Y. Cheng, “Faasnet: Scalable and fast provisioning of custom serverless container runtimes at alibaba cloud function compute,” in *2021 USENIX Annual Technical Conference, USENIX ATC 2021, July 14-16, 2021*. USENIX Association, 2021, pp. 443–457. [Online]. Available: <https://www.usenix.org/conference/atc21/presentation/wang-ao>
- [34] S. Fouladi, F. Romero, D. Iter, Q. Li, S. Chatterjee, C. Kozyrakis, M. Zaharia, and K. Winstein, “From laptop to lambda: Outsourcing everyday jobs to thousands of transient functional containers,” in *2019 USENIX Annual Technical Conference, ATC Renton, WA, USA, July 10-12, 2019*. USENIX Association, 2019, pp. 475–488. [Online]. Available: <https://www.usenix.org/conference/atc19/presentation/fouladi>
- [35] K. Kaffes, N. J. Yadwadkar, and C. Kozyrakis, “Hermod: principled and practical scheduling for serverless functions,” in *Proceedings of the 13th Symposium on Cloud Computing, SoCC 2022, San Francisco, California, November 7-11, 2022*. ACM, 2022, pp. 289–305. [Online]. Available: <https://doi.org/10.1145/3542929.3563468>
- [36] P. C. Castro, V. Ishakian, V. Muthusamy, and A. Slominski, “The rise of serverless computing,” *Commun. ACM*, vol. 62, no. 12, pp. 44–54, 2019. [Online]. Available: <https://doi.org/10.1145/3368454>
- [37] G. E. Noether, “Note on the kolmogorov statistic in the discrete case,” *Metrika*, vol. 7, pp. 115–116, 1963. [Online]. Available: <https://api.semanticscholar.org/CorpusID:120687545>
- [38] Microsoft. (2023) Azure functions triggers and bindings concepts. [Online]. Available: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-triggers-bindings>
- [39] —. (2023) Azure functions scenarios. [Online]. Available: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-scenarios?pivot=programming-language-csharp>
- [40] M. Azure and M. Research. (2019) Azure functions trace 2019. Microsoft Research. [Online]. Available: https://azurecloudpublicdataset2.blob.core.windows.net/azurepublicdatasetv2/azurefunctions_dataset2019/
- [41] A. Kuster, J. Shen, M. Dai, and J. Guo. (2022) reproducing-serverless-in-the-wild. [Online]. Available: <https://github.com/andreaskuster/reproducing-serverless-in-the-wild>
- [42] J. Wen, Z. Chen, Y. Liu, Y. Lou, Y. Ma, G. Huang, X. Jin, and X. Liu, “An empirical study on challenges of application development in serverless computing,” in *ESEC/FSE '21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Athens, Greece, August 23-28, 2021*. ACM, 2021, pp. 416–428. [Online]. Available: <https://doi.org/10.1145/3468264.3468558>
- [43] A. Tariq, A. Pahl, S. Nimmagadda, E. Rozner, and S. Lanka, “Sequoia: enabling quality-of-service in serverless computing,” in *SoCC '20: ACM Symposium on Cloud Computing, Virtual Event, USA, October 19-21, 2020*. ACM, 2020, pp. 311–327. [Online]. Available: <https://doi.org/10.1145/3419111.3421306>
- [44] F. Romero, G. I. Chaudhry, I. n. Goiri, P. Gopa, P. Batum, N. J. Yadwadkar, R. Fonseca, C. Kozyrakis, and R. Bianchini, “FaaS: A transparent auto-scaling cache for serverless applications,” in *Proceedings of the ACM Symposium on Cloud Computing*, ser. SoCC '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 122–137. [Online]. Available: <https://doi.org/10.1145/3472883.3486974>
- [45] H. Yu, A. A. Irissappane, H. Wang, and W. J. Lloyd, “Faasrank: Learning to schedule functions in serverless platforms,” in *IEEE International Conference on Autonomic Computing and Self-Organizing Systems, ACSOS 2021, Washington, DC, USA, September 27 - Oct. 1, 2021*. IEEE, 2021, pp. 31–40. [Online]. Available: <https://doi.org/10.1109/ACSOS52086.2021.00023>
- [46] Y. Zhang, I. Goiri, G. I. Chaudhry, R. Fonseca, S. Elnikety, C. Delimitrou, and R. Bianchini, “Faster and cheaper serverless computing on harvested resources,” in *SOSP '21: ACM SIGOPS 28th Symposium on Operating Systems Principles, Virtual Event / Koblenz, Germany, October 26-29, 2021*. ACM, 2021, pp. 724–739. [Online]. Available: <https://doi.org/10.1145/3477132.3483580>
- [47] R. B. Roy, T. Patel, and D. Tiwari, “Icebreaker: warming serverless functions better with heterogeneity,” in *ASPLOS '22: 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Lausanne, Switzerland, 28 February 2022 - 4 March 2022*. ACM, 2022, pp. 753–767. [Online]. Available: <https://doi.org/10.1145/3503222.3507750>
- [48] N. Daw, U. Bellur, and P. Kulkarni, “Xanadu: Mitigating cascading cold starts in serverless function chain deployments,” in *Middleware '20: 21st International Middleware Conference, Delft, The Netherlands, December 7-11, 2020*. ACM, 2020, pp. 356–370. [Online]. Available: <https://doi.org/10.1145/3423211.3425690>
- [49] B. Sethi, S. K. Addya, and S. K. Ghosh, “Lcs: Alleviating total cold start latency in serverless applications with lru warm container approach,” in *Proceedings of the 24th International Conference on Distributed Computing and Networking*. New York, NY, USA: Association for Computing Machinery, 2023, p. 197–206. [Online]. Available: <https://doi.org/10.1145/3571306.3571404>
- [50] P. Silva, D. Fireman, and T. E. Pereira, “Prebaking functions to warm the serverless cold start,” in *Middleware '20: 21st International Middleware Conference, Delft, The Netherlands, December 7-11, 2020*. ACM, 2020, pp. 1–13. [Online]. Available: <https://doi.org/10.1145/3423211.3425682>
- [51] 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 2018, Boston, MA, USA, July 11-13, 2018*. USENIX Association, 2018, pp. 923–935. [Online]. Available: <https://www.usenix.org/conference/atc18/presentation/akkus>
- [52] Z. Li, L. Guo, Q. Chen, J. Cheng, C. Xu, D. Zeng, Z. Song, T. Ma, Y. Yang, C. Li, and M. Guo, “Help rather than recycle: Alleviating cold startup in serverless computing through inter-function container sharing,” in *2022 USENIX Annual Technical Conference, ATC Carlsbad, CA, USA, July 11-13, 2022*. USENIX Association, 2022, pp. 69–84. [Online]. Available: <https://www.usenix.org/conference/atc22/presentation/li-zijun-help>
- [53] K. Suo, Y. Shi, X. Xu, D. Cheng, and W. Chen, “Tackling cold start in serverless computing with container runtime reusing,” in *Proceedings of the 2020 Workshop on Network Application Integration/CoDesign, NAI@SIGCOMM 2020, Virtual Event, USA, August 14, 2020*. ACM, 2020, pp. 54–55. [Online]. Available: <https://doi.org/10.1145/3405672.3409493>
- [54] A. Mohan, H. S. Sane, K. Doshi, S. Edupuganti, N. Nayak, and V. Sukhomlinov, “Agile cold starts for scalable serverless,” in *11th USENIX Workshop on Hot Topics in Cloud Computing, HotCloud 2019, Renton, WA, USA, July 8, 2019*. USENIX Association, 2019. [Online]. Available: <https://www.usenix.org/conference/hotcloud19/presentation/mohan>