

Affinity-aware Serverless Function Scheduling

Giuseppe De Palma^{*†}, Saverio Giallorenzo^{*†}, Jacopo Mauro[‡], Matteo Trentin^{*†‡}, Gianluigi Zavattaro^{*†}

^{*}Università di Bologna
Bologna, Italy

[†]INRIA
Sophia Antipolis, France

[‡]University of Southern Denmark
Odense, Denmark

{giuseppe.depalma2, saverio.giallorenzo2, matteo.trentin2, gianluigi.zavattaro}@unibo.it, mauro@imada.sdu.dk

Abstract—Functions-as-a-Service (FaaS) is a Serverless Cloud paradigm where a platform manages the scheduling (e.g., resource allocation, runtime environments) of stateless functions. Recent work proposed using domain-specific languages to express per-function policies, e.g., policies that enforce the allocation on nodes that enjoy lower latencies to databases and services used by the function. Here, we focus on *affinity-aware* scenarios, i.e., where, for performance and functional requirements, the allocation of a function depends on the presence/absence of other functions on nodes.

We present **aAPP**, an extension of a declarative, platform-agnostic language that captures affinity-aware scheduling at the FaaS level. We implement an **aAPP**-based prototype on Apache OpenWhisk. Besides proving that a FaaS platform can capture affinity awareness using **aAPP** and improve performance in affinity-aware scenarios, we use our prototype to show that **aAPP** imposes no noticeable overhead in scenarios without affinity constraints.

Index Terms—Serverless, Function-as-a-Service, Function Scheduling, Affinity-awareness

I. INTRODUCTION

Functions-as-a-Service (FaaS) is a programming paradigm supported by the Serverless Cloud execution model [1]. In FaaS, developers implement a distributed architecture from the composition of stateless functions and delegate concerns like execution runtimes and resource allocation to the serverless platform, thus focusing on writing code that implements business logic rather than worrying about infrastructure management. The main cloud providers offer FaaS [2]–[4] and open-source alternatives exist too [5]–[8].

A common denominator of these platforms is that they manage the allocation of functions over the available computing resources, also called *workers*, following opinionated policies that favour some performance principle. Indeed, effects like *code locality* [7]—due to latencies in loading function code and runtimes—or *session locality* [7]—due to the need to authenticate and open new sessions to interact with other services—can substantially increase the run time of functions. The breadth of the design space of serverless scheduling policies is witnessed by the growing literature focused on

techniques that mix one or more of these locality principles to increase the performance of function execution, assuming some locality-bound traits of functions [9]–[14]. Besides performance, functions can have functional requirements that the scheduler shall consider. For example, users might want to ward off allocating their functions alongside “untrusted” ones—common threat vectors in serverless are limited function isolation and the ability of functions to (surreptitiously) gather weaponisable information on the runtime, the infrastructure, and the other tenants [15]–[18].

Although one can mix different principles to expand the profile coverage of a given platform-wide scheduler policy, the latter hardly suits all kinds of scenarios. This shortcoming motivated De Palma et al. [19], [20] to introduce a domain-specific, platform-agnostic, declarative language, called *Allocation Priority Policies* (APP) to specify custom function allocation policies. Thanks to APP, the same platform can support different scheduling policies, each tailored to meet the specific needs of a set of related functions. De Palma et al. validated their approach by implementing an APP-based serverless platform as an extension of the open-source Apache OpenWhisk project [19], [21].

Our contributions originate from the observation that, at lower levels of the cloud stack, popular Infrastructure-as-a-Service (IaaS) platforms (e.g., OpenStack [22]) and Container-as-a-Service (CaaS) systems (e.g., Kubernetes [23]) allow users to express affinity and anti-affinity constraints about the allocation of VM/containers—e.g., reducing overhead, by shortening data paths via co-location (affinity), increasing reliability, by evenly distributing VM/containers among different nodes (anti-affinity) or improving security, by preventing the co-location of VM/containers belonging to different trust tiers (anti-affinity). On the contrary, FaaS platforms do not natively support the possibility to express affinity-aware scheduling, where function allocation depends on the presence (affinity) or absence (anti-affinity) at scheduling time of other functions in execution on the available workers.

Contribution: Recognising the potential of FaaS-level affinity-aware scheduling policies, we propose a language-based solution, obtained by extending APP into an affinity-aware function scheduling language called **aAPP**. In Sec. II, we present an example of affinity-aware scheduling at the FaaS level that we use to informally introduce **aAPP**. We formalise our proposal in Sec. III, where we present the **aAPP** syntax

Research partly supported by project PNRR CN HPC - SPOKE 9 - Innovation Grant LEONARDO - TASI - RTMER funded by the NextGenerationEU European initiative through the MUR, Italy (CUP: J33C22001170001), the research project FREEDA (CUP: I53D23003550006) funded by the framework PRIN 2022 (MUR, Italy), the French ANR project SmartCloud ANR-23-CE25-0012, and the Independent Research Fund Denmark, grant no. 4283-00007B.




Fig. 1. Examples of a FaaS infrastructure (left) and an aAPP script (right).

and discuss the increment of expressiveness w.r.t. APP. In Sec. IV, we concretise our proposal by presenting a prototype implementation of an aAPP-based serverless platform, namely, an extension of Apache OpenWhisk able to enforce aAPP-defined FaaS (anti-)affinity scheduling constraints. In Sec. V, we experimentally show that the usage of (anti-)affinity constraints are beneficial by considering an implementation of the affinity-aware scenario introduced in Sec. II. In Sec. VI, we compare the performance of our aAPP-based prototype and vanilla OpenWhisk with 7 benchmarks to show that aAPP imposes negligible overhead. We discuss related work and draw concluding remarks in Sec. VII.

II. EXAMPLE OF AN AFFINITY-AWARE FAAS SCENARIO

We have a *divide-et-impera* data-crunching serverless application implemented through two companion functions. The first, invoked by the users, is called *divide*. Its task is to split some data into chunks, store them in a database, and invoke instances of the second function. The second function, invoked by the *divide* for each stored chunk, is called *impera*. Its task is to retrieve and process a chunk of data from the database.

We run the above functions on the FaaS infrastructure depicted on the left of Fig. 1. The infrastructure includes two zones (e.g., separate regions of a cloud provider) and it has a *Gateway* that decides on which worker to allocate the execution of the functions. The infrastructure also includes three workers: w_1 and w_2 in *Zone₁* and w_3 in *Zone₂*. Each zone hosts an instance of an eventually-consistent distributed database [24], used by the functions running in that zone—eventually-consistent systems are typical for (FaaS) scenarios like ours, where one favours throughput and availability w.r.t. e.g., overall data consistency [25].

In Fig. 1, we represent function allocation requests with labelled document icons sent to the *Gateway*. Note that the users (the laptop icons in Fig. 1) launch the *divide* function (e.g., d_3) while the running *divide* (e.g., d_2 requesting i_2 and i'_2) invoke the *impera* functions.

Our FaaS infrastructure executes other functions besides the one above. In Fig. 1, we represent these requests with the labels h_1 , h_2 , and h_3 which are compute-intensive functions—

called *heavy*—that use a high amount of computational resources of the worker running them.

Given this context, an initial example of an affinity-aware scheduling policy is to avoid the co-occurrence of the *divide* and *impera* functions with the *heavy* ones. In this way, we can improve the performance of *divide* and *impera* by avoiding resource contention with the *heavy* functions. Another improvement regards the interaction with the database. The eventually-consistent behaviour of the database entails possible delays to synchronise the instances. Waiting for synchronisation is necessary only when the functions accessing the database connect to different database instances. Moreover, to further reduce delay, we can exploit the principle of *session locality* and let functions running on the same worker share the same connection with the database. This affinity-aware scheduling policy places *impera* functions only on workers that already host *divide* functions and avoid the overhead of re-establishing new connections.

These constraints can be encoded in aAPP as shown in the script in Fig. 1. In the code, we find three top-level items: *d*, *i*, and *h*, which are tags that identify policies, each describing the scheduling logic of a set of related functions. In the example, the tag *d* describes the logic for the *divide* functions while *i* and *h* target respectively the *impera* and *heavy* ones. The line `workers: *` found under all tags indicates that their related functions can use any of the available workers. From the top, under tag *d*, we use the `affinity` clause, introduced by aAPP, to specify that *d*-tagged functions should *not* be scheduled on a worker that currently hosts *heavy* functions (`!h`). Specifically, this is an example of *anti-affinity*, where we prevent the allocation of the tagged functions (e.g., *d*) on a worker that already hosts any anti-affine function (e.g., tagged *h*). Tag *i* declares the same anti-affinity for *heavy* functions, but it also indicates that *i*-tagged functions are *affine* with *d*-tagged ones. Affinity means that we can schedule a function on a candidate worker only if it currently hosts the former's affine functions. In the example, we use affinity to have *impera* functions run in the same worker of *divide* functions. Finally, we use tag *h* to complement the anti-affinity relation expressed in the previous tags, i.e., the *heavy* functions are anti-affine

```

                                id ∈ Identifiers      n ∈ ℕ

app ::=  $\overline{-tag}$ 
tag ::= id :  $\overline{-block}$  followup : f_opt
block ::= workers : w_opt strategy : s_opt
                                invalidate :  $\overline{-i_opt}$  affinity :  $\overline{-a_opt}$ 
w_opt ::= * |  $\overline{-id}$ 
s_opt ::= any | best_first
i_opt ::= capacity_used n% | max_concurrent_invocations n
a_opt ::= id | !id
f_opt ::= default | fail

```

Fig. 2. aAPP syntax.

```

- f_tag:
- workers:
  - local_w1
  - local_w2
  strategy: best_first
  invalidate:
  - capacity_used 80%
  affinity: g_tag,!
  h_tag
- workers:
  - public_w1
followup: fail

```

Fig. 3. Example aAPP script.

with both *d* and *i* functions and shall not be scheduled in workers that already host any of the latter.

Notably, we purposefully do not identify who writes the aAPP script in the example, e.g., the developer of the functions or the administrator of the platform. Indeed, aAPP (in general, APP and all its extensions) caters to different cloud stakeholders for scheduling policy definition. For instance, if we contextualise our example in a local private cloud setup, then users can directly write their own aAPP scripts because they have direct knowledge of the infrastructure nodes. Contrarily, if we are in a managed cloud environment, the cloud provider would use aAPP to implement and enforce scheduling requirements specified by their clients based on their workflows—e.g., synthesising aAPP scripts from function and workflow code [26], [27].

III. THE AAPP LANGUAGE

We now present aAPP, our extension of the FaaS function scheduling language APP [19], [20] with affinity and anti-affinity constraints.

We report in Fig. 2 the syntax of aAPP. From here on, we indicate syntactic units in *italics*, optional fragments in *grey*, terminals in monospace, and lists with *bars*. The syntax of aAPP draws inspiration from YAML [28], a renowned data-serialisation language for configuration files—e.g., many modern cloud tools, like Kubernetes and Ansible, use this format.¹ In aAPP, functions have associated a tag that identifies some scheduling policies. An aAPP script represents: *i*) named scheduling policies identified by a *tag* and *ii*) policy *blocks* that indicate either some collection of workers, each identified by a worker *id*, or the universal *. To schedule a function, we use its tag to retrieve the scheduling policy that includes one or more blocks of possible workers. To select the worker, we iterate top-to-bottom on the blocks. We stop at the first block that has a non-empty list of valid

¹While aAPP scripts are YAML-compliant, for presentation, we stylise the syntax to increase readability. For instance, we omit quotes around strings, e.g., * instead of "*".

workers and then select one of those workers according to the strategy defined by the block (described later).

Each tag can define a *followup* clause, which specifies what to do if the policy of the tag did not lead to the scheduling of the function; either *fail*, to terminate the scheduling, or *default*, to apply the special default-tagged policy. Each block can define a *strategy* for worker selection (*any* selects non-deterministically one of the available workers in the list; *best_first* selects the first available worker in the list), a list of constraints that *invalidates* a worker for the allocation (*capacity_used* invalidates a worker if its resource occupation reaches the set threshold; *max_concurrent_invocations* invalidates a worker if it hosts more than the specified number of functions), and an *affinity* clause that carries a list containing affine tag identifiers *id* and anti-affine tags, represented by negated tag identifiers *!id*. aAPP is a minimal extension of APP where we add the *affinity* clause to capture (anti-)affinity constraints. Similar to the notion of affinity introduced by Microsoft in its IaaS offering [29], in aAPP, the relation of (anti-)affinity is “directional”, i.e., we impose no well-formedness property like symmetry or anti-symmetry on affinity or anti-affinity to avoid limiting aAPP’s ability to capture meaningful scenarios.²

We practically illustrate aAPP with the two-block aAPP policy example (for functions tagged *f_tag*) found in Fig. 3. The first block restricts allocations on the workers labelled *local_w1* and *local_w2* and the latter on *public_w1*. The first block specifies as invalid (i.e., such that they cannot host the function under scheduling) the workers that reach a memory consumption above 80%. Since the strategy is *best_first*, we allocate the function on the first valid worker; if none are valid, we proceed with the next block. The function has affinity with *g_tag* and anti-affinity with *h_tag*. Hence, a valid

²For instance, if we had symmetric anti-affinity, we would not capture scenarios where, e.g., a function *init* is the seeding function for a database and function *query* manipulates that data. The function *init* should always run before *query* but never where *query* is already running, while function *query* should run where *init* is present. To obtain this behaviour, we need *init* anti-affine with *query* but *query* affine with *init*.

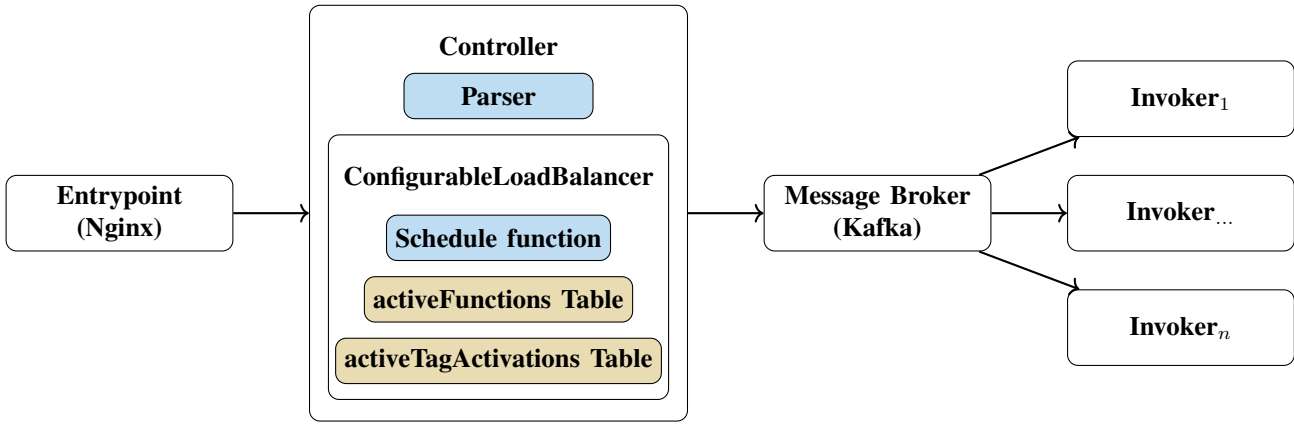


Fig. 4. Extended Apache OpenWhisk for aAPP (modified modules in blue, added modules in yellow).

worker must host at least a function with tag `g_tag` and no functions with tag `h_tag`. If both the first and second blocks do not find a valid worker, the scheduling of the function **fails** (instead of continuing with the default tag).

Notably, the addition of (anti-)affinity constraints increases the expressiveness of APP. Indeed, APP can capture anti-affinity constraints by severely limiting the flexibility of resource allocation, e.g., either through a partition of the workers and a static assignment of anti-affine functions to distinct partitions, or by limiting artificially the capacity that can be used or the number of functions we can allocate on a worker. This approach conflicts with the cloud principle of resource sharing and optimization. The situation for affinity is even poorer: APP cannot capture these constraints because it does not keep track of the functions allocated on workers.

IV. AAPP-BASED APACHE OPENWHISK

To validate aAPP, we implemented an aAPP-based FaaS platform, obtained by the existing APP-based version of Apache OpenWhisk.

We support the description of our implementation with Fig. 4, which illustrates the main components and the typical flow of function invocations in OpenWhisk. Specifically, the *Entrypoint* for function execution requests is an Nginx reverse proxy, which forwards the requests to a *Controller*. The *Controller*, in turn, routes them via a Kafka *Message Broker* to the *Invokers*—the workers in OpenWhisk terminology.

In the figure, we highlight the main interventions we performed to make the existing APP-based OpenWhisk version aAPP-compliant. The extension concerns two parts: the APP parser and the *ConfigurableLoadBalancer*, both absent in vanilla OpenWhisk and originally introduced in APP-based OpenWhisk [19]. We respectively extended the parser and the *ConfigurableLoadBalancer* to add compatibility for aAPP scripts and to keep track of the functions allocated to all the workers. In particular, we introduced two lookup tables, called *activeFunctions* and *activeTagActivations*, to implement the tracking functionality. The first table associates the allocated functions (and their tags) to their host worker and allows

the load balancer to verify (anti-)affinity constraints. The *activeTagActivations* table keeps tracks of the state of the different function instances (possibly of the same function definition, so we cannot use their identifiers) by pairing their *activation ids* with their function identifiers. When we observe the termination of an active function, we look its function identifier up and remove that instance from the *activeFunctions* table—we detect instance terminations thanks to the messages workers send to notify the load balancer of their completion. Code-wise, these changes required the modification of the *schedule* function of the *ConfigurableLoadBalancer*.

The scheduling semantics of aAPP scripts is straightforward. We present it in (Python-like) pseudo-code in Listing 1. In Listing 1, the *schedule* function requires the name of the function under scheduling (`f`), a map that represents the current infrastructure configuration (workers and functions running therein, explained later) (`conf`), an aAPP script encoded as a Python dictionary of objects (`aapp`), and a registry that maps each function to a tag and its memory occupancy (`reg`). The infrastructure configuration maps, for each worker, the list of functions scheduled on it (`fs`), the memory allocated for those functions (`memory_used`), and the total amount of memory of the worker (`max_memory`).

Given these inputs, *schedule* gets the tag associated with `f` (Line 2) and then extracts the blocks associated with this tag in the `aapp` script (Line 3). If the `followup` option is different from “fail”, we append the blocks associated with the default tag to the list of `f`’s blocks (Line 5). Then, we obtain the list of valid workers for every block in order of appearance (Line 9). When the `workers` clause uses `*`, we consider all the workers present in the configuration (Line 8). If the list of valid workers is non-empty, we choose the first one when the strategy is `best_first` (Line 12) and a random one otherwise (Line 14). If the list is empty, the scheduling fails (Line 15).

The *schedule* function uses the valid function to check when a worker is valid, i.e., it is available, it has enough capacity to host the function (Lines 18–19), and that allocating on it the function satisfies all the constraints

```

1 def schedule(f, conf, aapp, reg):
2     (memory, tag) = reg[f]
3     blocks = aapp[tag].blocks # get the blocks
4     if aapp[tag].followup != 'fail':
5         blocks += aapp['default'].blocks # add default tag blocks
6     for block in blocks:
7         if '*' in block['workers']:
8             block['workers'] = conf.keys
9             workers = [ for worker in block['workers'] if valid(f,worker,conf,reg,block)]
10            if len(workers) > 0: # if at least one valid worker is found
11                if block['strategy'] == 'best_first':
12                    return workers[0]
13                elif block['strategy'] == 'any':
14                    return random.choice(workers)
15            raise Exception('Function not schedulable')

17 def valid(f, w, conf, reg, block):
18     (memory, tag) = reg[f]
19     if (w not in conf) or (conf[w]['memory_used'] + memory > conf[w]['max_memory']):
20         return False
21     if 'invalidate' in block:
22         if ('capacity_used' in block['invalidate']) and
23             (block['invalidate']['capacity_used'] <= conf[w]['memory_used']):
24             return False
25         if ('max_concurrent_invocations' in block['invalidate']) and
26             (block['invalidate']['max_concurrent_invocations'] <= len(conf[w]['fs'])):
27             return False
28     if 'affinity' in block:
29         affine_tags = set([t for t in block['affinity'] if not t.startswith('!')])
30         anti_affine_tags = set([t[1:] for t in block['affinity'] if t.startswith('!')])
31         w_tags = set([t for (_, t) in [reg(f) for f in conf[w]['fs']]])
32         for t in affine_tags:
33             if t not in w_tags: return False
34         for t in anti_affine_tags:
35             if t in w_tags: return False
36     return True

```

Listing 1. The pseudo-code of the schedule and valid functions.

of `capacity_used`, `max_concurrent_invocations` (Lines 21–26), and `affinity` (Lines 27–34).

To implement our prototype, we have modified the Scala codebase of the OpenWhisk project; specifically, we have modified the scheduling algorithm to implement the logic of Listing 1 and manage the workers-functions status, on a fork of the OpenWhisk repository [30]. The entire system is easily deployable using Terraform and Ansible scripts.

V. PERFORMANCE IMPROVEMENTS VIA AFFINITY-AWARENESS

To validate our platform and show the benefit of (anti-)affinity constraints for affinity-aware scenarios, we use the example from Sec. II as a benchmark case study and show that, by enforcing (anti-)affinity constraints, we can reduce average execution times and tail latency.

Recalling the example, we consider a simple *divide-et-impera* serverless architecture running in a realistic co-tenancy context. Users invoke *divide* functions, requesting the solution of a problem. At invocation, *divide* splits the problem into sub-problems and invokes instances of the second function, *impera*. The *impera* instances solve their relative sub-problems and store their solution fragments on a persistent storage service. After the *imperas* terminated, *divide* retrieves the partial solutions, assembles them, and returns the response to the user. We consider a multi-zone execution context where each zone hosts an instance of an eventually-consistent distributed database. The workers in one zone access the local instance of the database. Another function, called *heavy*, represents possible interferences of serverless co-tenancy.

Experimental setup. To run the case study, we deploy the OpenWhisk versions of APP and aAPP on an 8-node Kuber-

```

1 - d:
2   - workers: *
3     strategy: random
4     affinity:
5       - !h_eu
6       - !h_us
7
8 - i:
9   - workers: *
10  strategy: random
11  affinity:
12    - !h_eu
13    - !h_us
14    - d
15
16 - h_eu:
17   - workers:
18     workereu1
19 - h_us:
20   - workers:
21     workerus1
22

```

Fig. 5. The aAPP script used for the tests.

netes cluster on the Digital Ocean platform. As schematised in Fig. 7, one node acts as the control plane (and as such, it is unavailable to OpenWhisk), one hosts the OpenWhisk core components (i.e., the Controller, the OpenWhisk internal database CouchDB, and the messaging system Kafka), and six nodes are workers. We deploy the control plane and the OpenWhisk core components on virtual machines with 2 vCPU and 2 GB RAM, while we deploy 4 workers on virtual machines with 2 vCPU and 2 GB RAM and 2 workers with 1 vCPU and 1 GB RAM. All machines run the Ubuntu Server 20.04 OS. Location-wise, we place the control plane, the OpenWhisk core components, and 3 workers in Europe and 3 workers in North America (2 with the more powerful configuration and 1 with the lesser one in each zone). To implement persistent storage, we deploy a 2-node MongoDB replica set, one in Europe and one in North America, using the 6.0.2 version of the Community Server.

We distribute the load generated by the *heavy* functions on the platform with two variants, *heavy_eu* and *heavy_us*, which we constrain to be resp. allocated in the Europe and the North America data centres on the less powerful workers (identified with *workereu1* and *workerus1*), to further amplify the effect of co-tenancy they exert.

All functions are in JavaScript and run on OpenWhisk NodeJS runtime *nodejs:14*. The *divide* function invokes two instances of the *impera* functions with 3 parameters: i) a freshly generated index, ii) an array populated with 100 random numbers, and iii) the initial-final indexes of the values to work on (0–49 to the first *impera* instance, 50–99 to the second one). The *divide* function waits for the *impera* functions to terminate, and then opens a connection to the local instance of MongoDB with the aim of retrieving documents representing the results computed by the *impera* functions. The *divide*

function detects the correct documents by issuing a query on the MongoDB with the freshly generated index. Each *impera* function simulates a computation on the values received from the *divide* function by connecting to the local instance of MongoDB and storing on it one document for each received value. Each of these documents contains one of the values and the freshly generated index received upon invocation. In this way, each execution of the *impera* function stores in MongoDB 50 documents. The *heavy* function simulates computing-resource consumption by performing 1 billion iterations of a computation consisting of the random generation of two numbers and the execution of their multiplication.

Note that the *divide* function needs to retrieve from MongoDB the documents generated by the two *impera* function it invokes. Each function opens a connection to the local instance of MongoDB and, in case an *impera* function runs on a different zone w.r.t. the *divide* function, it can take some time for the two database instances to eventually become consistent (converge). The *divide* function implements an exponential back-off retry approach [31]—it tries to fetch the documents from its local storage instance; if the data is not there, starting from a 1-second delay, the function waits for a back-off time that exponentially increases at each retry.

Experiments and Results: In our experiment, we consider three APP/aAPP scripts to showcase the benefits of (anti-)affinity constraints. The first, which uses the full expressiveness of aAPP, is the one reported in Fig. 5—where *imperas* (tagged with *i*) are affine with *divide* (tagged with *d*) and they are both anti-affine with the *heavy* functions (tagged with *h_eu* and *h_us*). We impose affinity between *divide* and *impera* functions to guarantee that the database writer (the *impera* function) and reader (the *divide* function) access the same database instance. We impose anti-affinity with the *heavy* functions to avoid resource contention between the *divide/impera* functions and the *heavy* functions. The second script removes the affinity constraints between *impera* and *divide* from the first script (anti-affinity-only-aAPP). The third script omits the anti-affinity constraints from the second one, effectively making it an APP script.

Each experiment involves 5 sequential runs. Each run invokes the *heavy_eu* and *heavy_us* functions in non-blocking mode, followed by 10 calls of the *divide* function, each one waiting for the previous to complete. Upon termination of the *heavy* functions, we proceed with the remaining runs, for a total of 10 *heavy* and 50 *divide* functions per experiment. To ensure reliable results, we run the experiment 5 times, totalling 250 calls of the *divide* function for each of the three APP/aAPP script. We use Apache JMeter to simulate each request, tracking its latency, number of retries (to retrieve storage data), and outcomes (success or failure).

We report the mean, median, and 95th tail latencies for the *divide* functions at the top of Fig. 6 under aAPP, Anti-Affinity-Only aAPP, and APP. From the values, we mainly note that aAPP has the best performances, while the latencies increase for both anti-affinity-only aAPP and, even more substantially, for APP.

Configuration	Mean Latency (ms)	Median Latency (ms)	95 th Tail Latency (ms)
aAPP	1547	883	3041
anti-affinity-only aAPP	2337 (+40%)	2381 (+91%)	3476 (+13%)
APP	8118 (+135%)	2648 (+99%)	60157 (+180%)




Fig. 6. Top, table reporting the mean, median, and 95th tail latencies of the *divide* functions under aAPP, anti-affinity-only aAPP, and APP (percentages represent variation w.r.t. aAPP). Bottom, the sorted scatter plot of the latencies incrementally sorted.

Role	vCPU	RAM	Location
Control Plane	2	2 GB	Europe
OW Core	2	2 GB	Europe
OW Worker	2	2 GB	Europe
OW Worker	2	2 GB	Europe
OW Worker	2	2 GB	North America
OW Worker	2	2 GB	North America
OW Worker	1	1 GB	Europe
OW Worker	1	1 GB	North America
MongoDB Replica	1	1 GB	Europe
MongoDB Replica	1	1 GB	North America

Fig. 7. Nodes used for the experimental evaluation.

To further analyse the differences between the performance of the three policies, at the bottom of Fig. 6, we report the sorted scatter plot of the latencies of the *divide* functions from the shortest to the longest (x -axis). We focus on this measure because it offers a comprehensive overview of the performance of the architecture. In particular, this measure includes the latencies of the related *impera* functions, whose run times concretely impact the request-response delay experienced by the users interacting with the system.

The first striking observation we gather is that there is a gap (there are almost no instances) in the distribution of the aAPP data points between the 1000ms and the 2400ms mark. We conjecture that this behaviour derives from having OpenWhisk core components installed in one region, which exert some

overhead on the workers of the other region when they interact with the platform (e.g., to fetch functions and receive/send requests/notifications). We see similar intervals, although less apparent, for APP and anti-affinity-only-aAPP.

In the 200–1000ms interval, aAPP provides consistent, fast performance, while APP and anti-affinity-only aAPP show only a few well-performing cases—the rest of the data points at the corresponding performance bracket are shifted to the right, revealing comparatively slower results. We can characterise the “fast” invocations as those where the *divide* and its two *impera* functions appear on a “free” node, i.e., without the *heavy* function, in Europe. Specifically, when using APP, each invocation has a $2/6$ probability of appearing on a free node in Europe, i.e., the probability of fast invocations is $(2/6)^3 \approx 3.7\%$, with anti-affinity-only aAPP the figure increases to $(1/2)^3 = 12.5\%$ (each invocation has a $1/2$ chance of appearing on a European free node) and with aAPP the probability raises to 50% since all three functions go on the same node (either in the US or in the EU).

Overall, already introducing anti-affinities improves performance (mean, median, tail latency improve resp. of 110%, 10%, and 178%), which shows the impact of sharing a worker with *heavy* functions—APP shows a long tail of invocations after the ca. 3000ms mark in Fig. 6. Looking at worst cases, using aAPP does not result in a considerable performance increase. This is visible from the plot by noticing how the tail high-percentage instances of anti-affinity-only aAPP and aAPP almost overlap, resulting in a small (+13%) improve-

Fig. 8. Comparison of scheduling times between vanilla, APP-, and aAPP-based OpenWhisk (avg and st dev are in ms).

	OpenWhisk		APP		aAPP	
	avg	st dev	avg	st dev	avg	st dev
<i>hello-world</i>	0.68	1.16	0.73	1.25	0.8	1.27
<i>long-running</i>	0.48	0.53	0.69	0.92	0.71	1.01
<i>compute-intens.</i>	11.57	11.92	10.17	11.67	10.01	9.66
<i>DB-acc., light</i>	0.65	1.31	0.85	1.62	0.83	1.31
<i>DB-acc., heavy</i>	0.44	0.69	0.91	1.25	1.04	1.7
<i>external service</i>	1.28	2.08	1.95	3.33	1.49	2.5
<i>code dependen.</i>	0.64	1.06	1.0	2.27	0.86	1.8

ment in tail latency. The differences in mean (+40%) and median (+91%) latency between having affinities or not emerge in the 250–1000ms bracket, where the lack of affinity leads to fewer fast executions, in contrast with the abundance of fast aAPP instances. Practically, the figures and distribution show how strongly North American allocations impact latency vs the benefit of co-location. Besides increasing performance, aAPP succeeds in eliminating database access retries, contrarily to anti-affinity-only aAPP (i.e., 42 requests suffer at least one retry in APP, 23 in anti-affinity-only aAPP, and 0 in aAPP).

VI. AAPP’S OVERHEAD IS NEGLIGIBLE

While aAPP allows us to exploit (anti-)affinity constraints that would not be expressible otherwise, it is crucial to assess the overhead introduced by the additional functionalities of aAPP. In this section, we show that the added functionalities (to track the state of functions on workers) of our aAPP-based prototype have negligible impact on the platform’s scheduling performance.

For these experiments, we use the benchmark suite used by De Palma et al. [21] to benchmark their APP-based OpenWhisk implementation. Note that, in the context of these experiments, we are not interested in the data locality capabilities of APP, but only in checking the scheduling performances of aAPP. Thus, we deploy the platforms in only one cloud zone and use 2000 invocations for each scenario, to simplify as much as possible the testing environment and have enough invocations to draw meaningful comparisons. The benchmarks are:

- *hello-world* implements a simple echo application, and indicates the baseline performance of the platform;
- *long-running* waits for 3 seconds before responding to benchmark the handling of multiple functions running for several seconds and the management of their queuing process;
- *compute-intensive* multiplies two 10^2 square matrices and returns the result to the caller, measuring both the performance of handling functions performing some meaningful computation and of handling large invocation payloads;
- *DB-access (light)* executes a query for a document from a remote MongoDB database. The requested document is lightweight, corresponding to a JSON document of 106 bytes, with little impact on computation. De Palma et

al. used the case to measure the impact of data locality on the overall latency. Since we have all workers in the same cloud zone, we use it to measure the overhead of scheduling functions that fetch small payloads from a local database;

- *DB-access (heavy)* regards both a memory- and bandwidth-heavy data-query function. The function fetches a large document (124.38 MB) from a MongoDB database and extracts a property from the returned JSON. Similarly to the previous function, we use this one to evaluate the overhead of scheduling functions that fetch large payloads from a local database;
- *External service* tracks the performance of invoking an external API (Slack). De Palma et al. drew the function from the Wonderless dataset [32];
- *Code dependencies* is a formatter that takes a JSON string and returns a plain-text one, translating the key-value pairings into Python-compatible dictionary assignments. De Palma et al. drew this case from the Wonderless dataset [32].

For completeness, we note that we omitted the *cold-start* case from De Palma et al. [21], which is an echo application with sizable, unused dependencies. The peculiarity of the case is its 10-minute invocation pattern, used to track cold-start times (so that the platform evicts cached copies of the function, requiring costly fetch-and-startup times at any subsequent invocation). We omit this test since we can observe its effects with the *hello-world* and *code-dependency* cases.

We run the benchmarks on a one-zone Google Cloud cluster with four Ubuntu 20.04 virtual machines with 4 GB RAM each, one with 2 vCPU for the OpenWhisk controller and three with 1 vCPU, resp. for two workers and a MongoDB instance for the *DB-access* cases. We run 2000 function invocations for each case in batches of 4 parallel requests (500 per thread), recording both the scheduling time (the time between the arrival of a request at the controller and the issuing of the allocation) and the execution latencies. We compare aAPP, APP, and vanilla OpenWhisk. For a fair comparison with vanilla OpenWhisk, we set the APP/aAPP configurations with a default policy that falls back to the vanilla scheduler.

For all cases and platforms, we report in Fig. 8, in tabular form, the average (avg) and standard deviation (st dev) of the scheduling time. On average, all platforms allocate functions in less than 2ms, except for the *compute-intensive* case, which takes less than 12ms (likely due to the large request payloads that the controller needs to forward to workers). As expected, OpenWhisk vanilla is the fastest, closely (under one millisecond) followed by APP and aAPP—except for the *compute-intensive* case, where APP and aAPP perform better and OpenWhisk is slower by less than 2ms. The differences between APP and aAPP are even smaller, with APP being generally slightly (sub-millisecond) faster than aAPP. To better characterise the comparison, in Fig. 9, we show the plot-line distribution of the scheduling times. The curves exhibit the typical tail distribution pattern [33] of cloud workloads (which accounts for the high standard deviation reported in Fig. 8) and

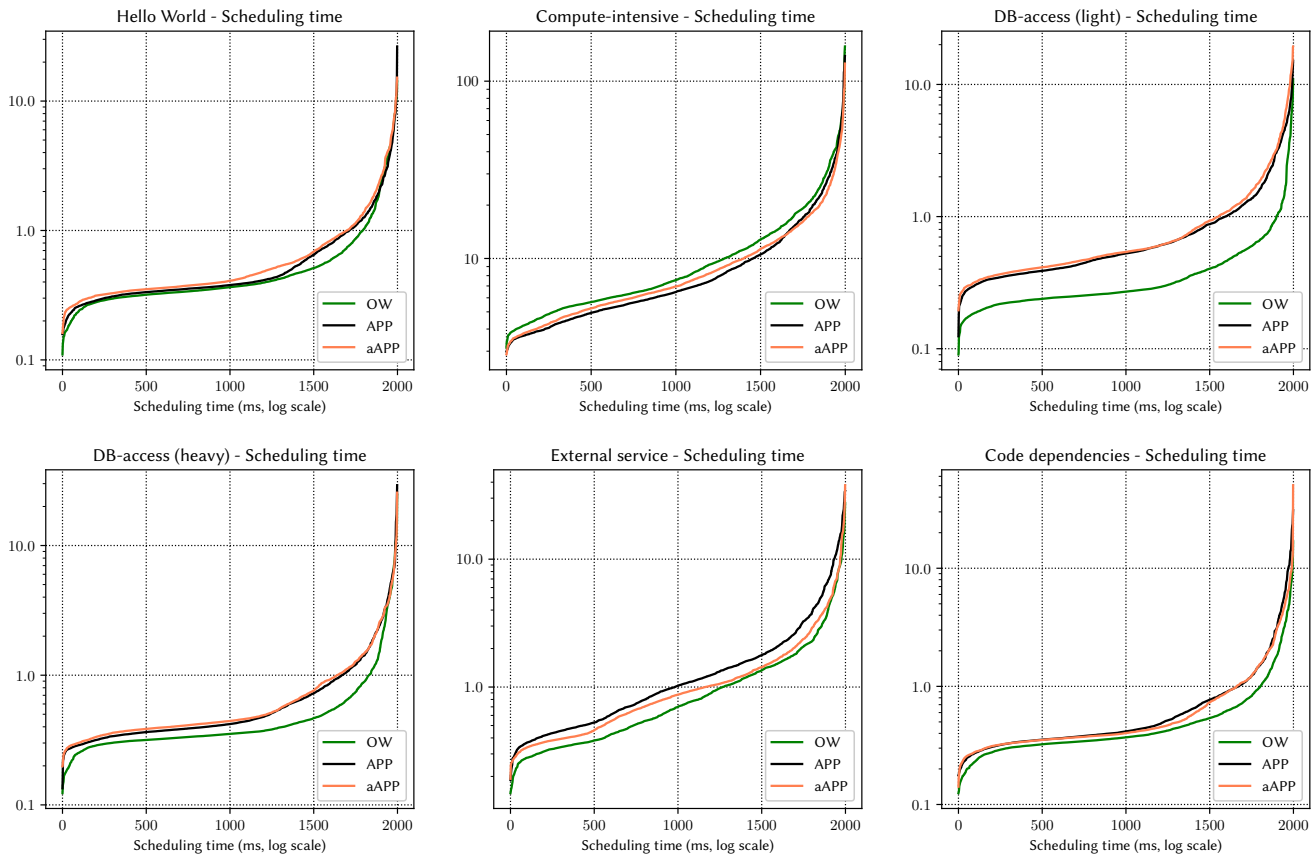


Fig. 9. Distribution of scheduling times for vanilla, APP-, and aAPP-based OpenWhisk (x-axis represents instances sorted from the quickest scheduling time to the slowest).

confirm our observations. Excluding the tails, the two plots almost overlap with negligible, sub-millisecond differences. All the results reported in Figs. 6 and 9 differ significantly according to the Wilcoxon test ($p = 0.001$).

VII. RELATED WORK AND CONCLUSION

To the best of our knowledge, aAPP is the first language that allows developers to state affinity constraints to better control the scheduling of the functions in FaaS platforms. By extending OpenWhisk, we demonstrate the effectiveness of using (anti-)affinity constraint of aAPP in reducing latency and tail latency. Furthermore, we benchmark that the overhead of supporting aAPP-based affinity constraints is minimal compared to vanilla OpenWhisk and its APP-based variant.

Broadening our scope, the works we see the closest to ours come from the neighbouring area of microservices [34]—the state-of-the-art style alternative to serverless for cloud architectures. Proposals in this direction are by Baarzi and George [35], who present a framework for the deployment of microservices that infers and assigns affinity and anti-affinity traits to microservices to orient the distribution of resources and microservices replicas on the available machines; Sampaio et al. [36], who introduce an adaptation mechanism for microservice deployment based on microservice affinities

(e.g., the more messages microservices exchange the more affine they are) and resource usage; Sheoran et al. [37], who propose an approach that computes procedural affinity of communication among microservices to make placement decisions. Looking at the industry, Azure Service Fabric [38] provides a notion of *service affinity* that ensures that the replicas of a service are placed on the same nodes as those of another, affine service. Another example is Kubernetes, which has a notion of *node affinity* and *inter-pod (anti-)affinity* to express advanced scheduling logic [23].

Overall, the mentioned work proves the usefulness of affinity-aware deployments at lower layers than FaaS (e.g., VMs, containers, microservices) and compels a discussion on the interplay between aAPP and IaaS/CaaS-level affinity, which we detail under two main directions. On the one hand, one could realise a version of aAPP for the Infrastructure and/or the Container layers. We argue it is more interesting to focus on FaaS. Indeed, there are mainstream IaaS and CaaS platforms that allow users to program directly ad-hoc schedulers (e.g., Kubernetes exposes APIs for creating scheduler plugins that define its scheduling policies). Since these layers afford a higher level of customisation than aAPP—at the expense of more technical involvement on the part of the users—a variant of aAPP for the IaaS/CaaS-levels seems less useful.

On the other hand, one can use IaaS and CaaS platforms that support affinity constraints to implement affinity-aware FaaS platforms. We see two main problems with pursuing this path. The first regards performance. To implement FaaS-level affinity using IaaS/CaaS affinity constraints, we need to impose a 1:1 relation between a function instance and the VM/container running it (if we let the same VM/container run parallel copies of the same function we cannot guarantee e.g., self anti-affinity). This imposition would prevent the platform from exploiting the serverless technique of VM/container reuse to avoid cold starts [39]–[41]. The second problem regards abstraction leakage, where letting FaaS users access the underlying IaaS/CaaS layers leaks details of the infrastructural components breaking FaaS’ paradigmatic abstractions.

A recent trend of FaaS is the definition/handling of the composition/workflows of functions, like AWS step-functions [42] and Azure Durable functions [43]. The main idea behind these works is to allow users to define workflows as the composition of functions with their branching logic, parallel execution, and error handling. The orchestrator/controller of the platform then uses the workflow to manage function executions and handle retries, timeouts, and errors. Our proposal is orthogonal to these works. Indeed, assuming the workflow is available, the orchestrator developed for handling serverless workflows should be extensible with an **aAPP**-like script to specify where to schedule the functions within a given workflow. Future work on this integration would support the enforcement of even more expressive policies than **aAPP**, like preventing function instances of the same workflow from sharing nodes.

Another interesting proposal, Palette [44], uses optional opaque parameters in function invocations to inform the load balancer of Azure Functions on the affinity with previous invocations and the data they produced. While Palette does not support (anti-)affinity constraints, it allows users to express which invocations benefit from running on the same node. We deem an interesting future work extending **aAPP** to infer (anti-)affinity constraints based on observed performance of different deployments or by inspecting the functions code. We already started exploring this landscape by introducing new options in **APP** that use static analysis and dynamic monitoring to select the workers that minimise the latency of calling external services in functions [26].

Regarding the constructs we have proposed for expressing affinity-aware policies in **aAPP**, we observe that an alternative approach could be to let the user directly declare the properties to enforce, leaving to the platform the task to realise them at run time. The scheduling runtime of this **APP** variant would allocate a function only if the allocation satisfies the formula or fail otherwise. The limitation of this approach lies in its scalability. Verifying the satisfiability of a property could require assessing multiple interacting constraints, possibly leading to an exponential time complexity with respect to the formula’s size, the number of workers, and the number of

functions involved.³ Contrarily, the **aAPP** scheduler checks whether it can allocate a function on a worker in linear time on the size of the workers and **aAPP** script length.

While, in this work, we focus on the design and run time performances of (anti-)affinity scheduling policies, we have also investigated the problem of statically checking scheduling properties of **APP** and **aAPP** scripts [46]. For instance, given the scheduling policy expressed by a script, one could be interested in checking the possibility for one safety critical function to be scheduled on an untrusted worker, or the possibility for the same function to be contemporaneously scheduled on the same worker with an “unknown” function developed by a third party. The approach we have investigated consists of the translation of the reachability property of interest into a planning problem, and then exploit off-the-shelf planners to solve such problem.

An important future research direction is empirically investigating the **aAPP**’s usability and intuitiveness by conducting experimental studies that probe developers’ ability to understand and specify (anti-)affinity constraints to implement properties of function scheduling (e.g., that two given functions are never scheduled on the same worker). Such investigations would both serve to validate the language’s design and uncover potential implicit understanding or unexpected interpretation patterns that might emerge when developers engage with **aAPP** with different levels of preconditioning.

Implementation-wise, we are planning to extend the support for **aAPP** to other serverless platforms. We already started this work by extending the support for a variant of **APP** in the FunLess serverless platform [47]. Natural steps in this direction include integrating **APP** with other popular FaaS platforms like Knative and OpenFaaS.

Regarding the integration within OpenWhisk, the platform supports scenarios where multiple controllers share the pool of available workers (e.g., for redundancy and load balancing) and take scheduling decisions without coordination. In our **aAPP**-based implementation, such multi-controller configurations present a problem since we need to prevent scheduling races among controllers—e.g., imagine two controllers that select an available, empty worker and, at the same time, allocate mutually anti-affine functions on it. Supporting multi-controller deployments is outside the scope of this paper and an interesting subject for future work.

Finally, while our evaluation demonstrates that affinity and anti-affinity constraints can enhance serverless application performance, a standardised benchmark of real-world examples would enable a more comprehensive analysis. Unfortunately, to the best of our knowledge, no such benchmark currently exists. As a direction for future work, we plan to collaborate with the community to collect instances of real-world applications and their affinity constraints, creating a benchmark to compare scheduler performances in serverless platforms.

³For example, if the language allows encoding properties such as “schedule the function f only if it does not prevent scheduling higher priority functions g_1, \dots, g_n ” then, due to the NP-hardness of bin-packing [45], the problem of scheduling a function becomes an NP-hard problem.

REFERENCES

- [1] E. Jonas, J. Schleier-Smith, V. Sreekanti, C.-C. Tsai, A. Khandelwal, Q. Pu, V. Shankar, J. Menezes Carreira, K. Krauth, N. Yadwadkar, J. Gonzalez, R. A. Popa, I. Stoica, and D. A. Patterson, "Cloud Programming Simplified: A Berkeley View on Serverless Computing," Tech. Rep. UCB/EECS-2019-3, University of California, 2019.
- [2] AWS, "AWS Lambda." <https://aws.amazon.com/lambda/>, 2024.
- [3] G. Cloud, "Google cloud functions." <https://cloud.google.com/functions/>, 2024.
- [4] Microsoft, "Microsoft azure functions." <https://azure.microsoft.com/>, 2024.
- [5] A. OpenWhisk, "Apache OpenWhisk." <https://openwhisk.apache.org/>, 2024.
- [6] OpenFaaS, "OpenFaaS." <https://www.openfaas.com/>, 2024.
- [7] S. Hendrickson, S. Sturdevant, E. Oakes, T. Harter, V. Venkataramani, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau, "Serverless Computation with OpenLambda," vol. 41, 2016.
- [8] Fission, "Fission." <https://fission.io/>, 2024.
- [9] G. Casale, M. Artač, W.-J. Van Den Heuvel, A. van Hoorn, P. Jakovits, F. Leymann, M. Long, V. Papanikolaou, D. Prezenza, A. Russo, *et al.*, "Radon: rational decomposition and orchestration for serverless computing," *SICS Software-Intensive Cyber-Physical Systems*, vol. 35, no. 1, pp. 77–87, 2020.
- [10] D. Kelly, F. Glavin, and E. Barrett, "Serverless computing: Behind the scenes of major platforms," in *2020 IEEE 13th International Conference on Cloud Computing (CLOUD)*, pp. 304–312, IEEE, 2020.
- [11] A. Banaei and M. Sharifi, "ETAS: predictive scheduling of functions on worker nodes of Apache OpenWhisk platform," *J. Supercomput.*, vol. 78, no. 4, pp. 5358–5393, 2022.
- [12] Z. Jia and E. Witchel, "Boki: Stateful Serverless Computing with Shared Logs," in *SOSP* (R. van Renesse and N. Zeldovich, eds.), pp. 691–707, ACM, 2021.
- [13] S. Kotni, A. Nayak, V. Ganapathy, and A. Basu, "Faastlane: Accelerating function-as-a-service workflows," in *Proc. of USENIX ATC*, pp. 805–820, USENIX Association, 2021.
- [14] C. P. Smith, A. Jindal, M. Chadha, M. Gerndt, and S. Benedict, "Fado: Faas functions and data orchestrator for multiple serverless edge-cloud clusters," in *ICFEC*, pp. 17–25, IEEE, 2022.
- [15] I. Baldini, P. Castro, K. Chang, P. Cheng, S. Fink, V. Ishakian, N. Mitchell, V. Muthusamy, R. Rabbah, A. Slominski, *et al.*, "Serverless computing: Current trends and open problems," in *Research advances in cloud computing*, pp. 1–20, Springer, 2017.
- [16] L. Wang, M. Li, Y. Zhang, T. Ristenpart, and M. Swift, "Peeking behind the curtains of serverless platforms," in *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, pp. 133–146, 2018.
- [17] K. Alpernas, C. Flanagan, S. Fouladi, L. Ryzhyk, M. Sagiv, T. Schmitz, and K. Winstein, "Secure serverless computing using dynamic information flow control," *Proc. ACM Program. Lang.*, vol. 2, no. OOPSLA, pp. 118:1–118:26, 2018.
- [18] P. Datta, P. Kumar, T. Morris, M. Grace, A. Rahmati, and A. Bates, "Valve: Securing function workflows on serverless computing platforms," in *Proceedings of The Web Conference 2020*, pp. 939–950, 2020.
- [19] G. De Palma, S. Giallorenzo, J. Mauro, and G. Zavattaro, "Allocation Priority Policies for Serverless Function-Execution Scheduling Optimization," in *ICSOC*, vol. 12571 of *Lecture Notes in Computer Science*, pp. 416–430, Springer, 2020.
- [20] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, and G. Zavattaro, "A declarative approach to topology-aware serverless function-execution scheduling," in *IEEE International Conference on Web Services, ICWS*, pp. 337–342, IEEE, 2022.
- [21] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, and G. Zavattaro, "An openwhisk extension for topology-aware allocation priority policies," in *Coordination Models and Languages - 26th IFIP WG 6.1 International Conference, COORDINATION 2024, Held as Part of the 19th International Federated Conference on Distributed Computing Techniques, DisCoTec 2024, Groningen, The Netherlands, June 17-21, 2024, Proceedings* (I. Castellani and F. Tiezzi, eds.), vol. 14676 of *Lecture Notes in Computer Science*, pp. 201–218, Springer, 2024.
- [22] OpenStack, "Documentation." <https://docs.openstack.org/project-deploy-guide/openstack-ansible/ocata/app-advanced-config-affinity.html>, 2024.
- [23] Kubernetes, "Node Affinity." <https://kubernetes.io/docs/tasks/configure-pod-container/assign-pods-nodes-using-node-affinity/>, 2024.
- [24] W. Vogels, "Eventually consistent," *Communications of the ACM*, vol. 52, no. 1, pp. 40–44, 2009.
- [25] P. Bailis and A. Ghodsi, "Eventual consistency today: Limitations, extensions, and beyond," *Communications of the ACM*, vol. 56, no. 5, pp. 55–63, 2013.
- [26] G. De Palma, S. Giallorenzo, C. Laneve, J. Mauro, M. Trentin, and G. Zavattaro, "Serverless scheduling policies based on cost analysis," in *Proceedings of the First Workshop on Trends in Configurable Systems Analysis, TiCSA@ETAPS 2023, Paris, France, 23rd April 2023* (M. H. ter Beek and C. Dubslaff, eds.), vol. 392 of *EPTCS*, pp. 40–52, 2023.
- [27] G. D. Palma *et al.*, "Towards a function-as-a-service choreographic programming language: Examples and applications," 2024.
- [28] YAML, "YAML specification." <https://yaml.org/spec/>, 11 2022.
- [29] Microsoft, "Service affinity in Service Fabric." <https://learn.microsoft.com/en-us/azure/service-fabric/service-fabric-cluster-resource-manager-advanced-placement-rules-affinity>, 2024.
- [30] G. D. Palma *et al.*, "<https://github.com/giusdp/openwhisk>," 2024.
- [31] P. Osman, *Microservices Development Cookbook: Design and build independently deployable, modular services*. Packt Publishing, 2018.
- [32] N. Eskandani and G. Salvaneschi, "The Wonderless Dataset for Serverless Computing," in *MSR*, pp. 565–569, IEEE, 2021.
- [33] J. Dean and L. A. Barroso, "The tail at scale," *Communications of the ACM*, vol. 56, no. 2, pp. 74–80, 2013.
- [34] N. Dragoni, S. Giallorenzo, A. Lluch-Lafuente, M. Mazzara, F. Montesi, R. Mustafin, and L. Safina, "Microservices: Yesterday, today, and tomorrow," in *Present and Ulterior Software Engineering*, pp. 195–216, Springer, 2017.
- [35] A. F. Baarzi and G. Kesidis, "Showar: Right-sizing and efficient scheduling of microservices," in *Proceedings of the ACM Symposium on Cloud Computing*, pp. 427–441, 2021.
- [36] A. R. Sampaio, J. Rubin, I. Beschastnikh, and N. S. Rosa, "Improving microservice-based applications with runtime placement adaptation," *Journal of Internet Services and Applications*, vol. 10, no. 1, pp. 1–30, 2019.
- [37] A. Sheoran, S. Fahmy, P. Sharma, and N. Modi, "Invenio: Communication affinity computation for low-latency microservices," in *Proceedings of the Symposium on Architectures for Networking and Communications Systems*, pp. 88–101, 2021.
- [38] Microsoft, "Azure service fabric." <https://learn.microsoft.com/en-us/azure/service-fabric/service-fabric-overview>, 2024.
- [39] A. Mohan, H. S. Sane, K. Doshi, S. Edupuganti, N. Nayak, and V. Sukhominov, "Agile Cold Starts for Scalable Serverless," in *Hot-Cloud*, USENIX Association, 2019.
- [40] V. Sreekanti, C. Wu, X. C. Lin, J. Schleier-Smith, J. E. Gonzalez, J. M. Hellerstein, and A. Tumanov, "Cloudburst: Stateful functions-as-a-service," *Proc. VLDB Endow.*, vol. 13, pp. 2438–2452, July 2020.
- [41] K. Solaiman and M. A. Adnan, "WLEC: A Not So Cold Architecture to Mitigate Cold Start Problem in Serverless Computing," in *IC2E*, pp. 144–153, IEEE, 2020.
- [42] AWS, "AWS Step Functions." <https://aws.amazon.com/step-functions/>, 2024.
- [43] S. Burckhardt, C. Gillum, D. Justo, K. Kallas, C. McMahon, and C. S. Meiklejohn, "Durable functions: semantics for stateful serverless," *Proceedings of the ACM on Programming Languages*, vol. 5, no. OOPSLA, pp. 1–27, 2021.
- [44] M. Abdi, S. Ginzburg, X. C. Lin, J. M. Faleiro, G. I. Chaudhry, I. Goiri, R. Bianchini, D. S. Berger, and R. Fonseca, "Palette Load Balancing: Locality Hints for Serverless Functions," in *EuroSys*, pp. 365–380, ACM, 2023.
- [45] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [46] G. De Palma, S. Giallorenzo, J. Mauro, M. Trentin, and G. Zavattaro, "Formally verifying function scheduling properties in serverless applications," *IT Prof.*, vol. 25, no. 6, pp. 94–99, 2023.
- [47] G. De Palma, S. Giallorenzo, C. Laneve, J. Mauro, M. Trentin, and G. Zavattaro, "Leveraging static analysis for cost-aware serverless scheduling policies," *International Journal on Software Tools for Technology Transfer*, Jan 2025.