



# Transparent serverless execution of Python multiprocessing applications

Aitor Arjona\*, Gerard Finol, Pedro García López

Universitat Rovira i Virgili, Tarragona, Spain



## ARTICLE INFO

### Article history:

Received 20 May 2022

Received in revised form 26 October 2022

Accepted 31 October 2022

Available online 8 November 2022

### Keywords:

Transparency

Access transparency

Serverless

FaaS

Multiprocessing

Parallel programming

## ABSTRACT

Access transparency means that both local and remote resources are accessed using identical operations. With transparency, unmodified single-machine applications could run over disaggregated compute, storage, and memory resources. Hiding the complexity of distributed systems through transparency would have great benefits, like scaling-out local-parallel scientific applications over flexible disaggregated resources in the Cloud.

This paper presents a performance evaluation where we assess the feasibility of access transparency over state-of-the-art Cloud disaggregated resources for Python multiprocessing applications. We have interfaced the multiprocessing module with an implementation that transparently runs processes on serverless functions and uses an in-memory data store for shared state.

To evaluate transparency, we run in the Cloud four unmodified applications: Uber Research's Evolution Strategies, Baselines-AI's Proximal Policy Optimization, Pandarallel's dataframe, and Scikit Learn's Hyperparameter tuning. We compare execution time and scalability of the same application running over disaggregated resources using our library, with the single-machine Python multiprocessing libraries in a large VM. For equal resources, applications efficiently using message-passing abstractions achieve comparable results despite the significant overheads of remote communication. Other shared-memory intensive applications do not perform due to high remote memory latency.

The results show that Python's multiprocessing library design is an enabler towards transparency: legacy applications using efficient disaggregated abstractions can transparently scale beyond VM limited resources for increased parallelism without changing the underlying code or architecture.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Coulouris et al. [1] define transparency as “the concealment from the user and application programmer of the complexities of distributed systems”. Access transparency allows to execute unmodified parallel code in a distributed environment, where resources (CPUs, memory) are distributed over many machines but accessed as if they were arranged in a single local machine.

The motivation of this paper is very simple: Low Latency  $\Rightarrow$  Disaggregation  $\Rightarrow$  Full Transparency. The downward trend in network latency [2,3] suggests that resource disaggregation is increasingly viable, which provides the opportunity to achieve access transparency in the next years [4–6]. Resource disaggregation in the Cloud has been the key to flexible scaling models provided by current serverless services such as Function-as-a-Service (FaaS) or Object Storage. Serverless services have also

proven to be effective for massive parallel computing applications [7,8] and even Big Data processing at scale [9].

If we can use identical operations for both local and remote resources with no significant performance degradation, it is then possible to unify the local and remote programming paradigms. In such a case, developer productivity would be greatly increased, as transparency would facilitate and make more accessible to program a distributed system, making the use of specific middleware redundant. Moreover, we could port legacy monolithic applications and scale them on a Cloud setting with flexible resources. New iterations of software modernization or architecture re-engineering could be avoided, saving maintenance costs and time to engineers. With transparency, resources could be adapted for compute-intensive applications to process larger workloads than a single machine could withstand without having to modify the underlying code or architecture.

As a matter of fact, in the European project Horizon 2020 Cloudbutton,<sup>1</sup> we aim to overly simplify and democratize the use

\* Corresponding author.

E-mail addresses: [aitor.arjona@urv.cat](mailto:aitor.arjona@urv.cat) (A. Arjona), [gerard.finol@urv.cat](mailto:gerard.finol@urv.cat) (G. Finol), [pedro.garcia@urv.cat](mailto:pedro.garcia@urv.cat) (P.G. López).

<sup>1</sup> <https://cloudbutton.eu>.

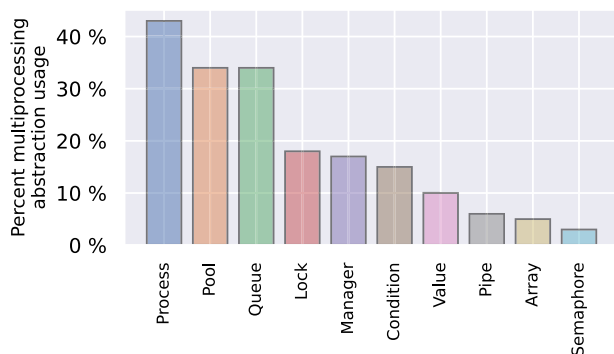


Fig. 1. Percentage of usage of the main multiprocessing abstractions of the top 100 most starred GitHub Python repositories that use multiprocessing.

of the Cloud for scientific computing using serverless technologies. Cloudbutton's core objective was inspired by a professor of computer graphics at UC Berkeley that wondered "Why is there no cloud button?" [7]. His students wished they could just "push a button" and have their existing single-machine code running on the Cloud. Therefore, transparency is an important consideration, because many of the tools currently used by data scientists are legacy applications that are difficult to parallelize or move to the Cloud. Moreover, data scientists are unlikely to be knowledgeable about Cloud technologies. Achieving access transparency could effortlessly unleash the potential of Cloud flexible and serverless resources to effectively speed up scientific computing pipelines.

Nonetheless, the distributed systems community has consistently criticized the idea of transparency. The reason is that Distributed Shared Memories (DSM), studied in depth in the past [10], suffered from complexity and performance issues, which made achieving transparency difficult. Waldo et al. [11] already discussed in 1994 that, in the context of Object Oriented Programming, the usage of remote objects as if they were local is incorrect and leads to performance issues and general unreliability. Although the authors did not explicitly discuss access transparency in their article, their conclusions suggest that remote memory access latency and partial failures make full transparency unfeasible.

The main criticism of transparency is that remote memory will never be as fast as local memory. Nevertheless, not all parallel programming models require intensive access to shared memory. On the contrary, many parallel applications just rely on communication and synchronization primitives that could be efficiently disaggregated.

In this line, the Python programming language uses processes to achieve true parallelism. Shared state in Python multiprocessing mainly consists of message passing (Pipes, Queues) or remote calls to other processes (Managers). To reinforce this point, we have analyzed the top 100 most starred repositories on GitHub that use Python multiprocessing and found out that Queues and Managers are the most used abstractions for state sharing (Fig. 1). In this aspect, the problems and limitations caused by DSMs would not apply, thus transparently adapting multiprocessing Python applications to a distributed environment is more feasible.

This article presents a performance study to evaluate if the inherent scalability of serverless functions, together with a disaggregated and consistent in-memory storage component, enables to transparently run unmodified Python multiprocessing applications over disaggregated serverless compute resources at scale. We want to run the same application with the same workload both in a VM (Virtual Machine) on AWS EC2 and with serverless functions on AWS Lambda, in order to compare execution time,

speedup, parallelism and to determine the possible overheads originated by moving to a distributed environment.

For this purpose, we have extended the Lithops serverless computing framework [12] with a module that fully implements the Python multiprocessing interface. This re-implementation leverages serverless functions for processes and Redis database for stateful multiprocessing abstractions (shared state, queues, locks...). Python applications written with the multiprocessing library can then be transparently ported to the Cloud by only changing the `import` statement.

For the performance study, we have used four scientific applications that make use of Python multiprocessing for parallelism: Evolution Strategies, Proximal Policy Optimization, Scikit-Learn Grid Search and Pandas dataframes. For each of them, we have compared its local execution on VMs currently available in AWS EC2 with its equivalent execution on AWS Lambda, while maintaining the same code, to assess if the application could be further scaled using serverless flexible resources despite the limitations and overheads generated.

The results show that it is actually possible to effectively run a local-parallel Python application using serverless resources, however, the application must efficiently use message-passing multiprocessing abstractions (Pipes, Queues, Managers...) in order to perform. On the contrary, applications that heavily use shared memory abstractions (Arrays) struggle to achieve competent results due to the high latency of remote memory. Yet, legacy applications that use efficiently disaggregated communication and synchronization abstractions yield good performance, opening the possibility to effortlessly scale on Cloud's flexible resources.

Our contributions are:

1. Performance study and evaluation of access transparency feasibility by comparing the execution of 4 unmodified applications (Evolution Strategies, Proximal Policy Optimization, Scikit-Learn Grid Search and Pandas data-frames) on a single VM and over serverless functions to analyze the speedup, parallelism and overheads.
2. We outline key insights from this study: (i) it is possible to run and scale unmodified Python multiprocessing applications with minor degradation compared to a VM, (ii) existing serverless overheads were partially masked by Hyper-Threading inefficiencies observed in a single VM, and (iii) Python's multiprocessing stateful abstractions clearly facilitated transparency (other languages directly accessing memory references or pointers would be more complex to intercept).
3. Extension to the Lithops computing framework with a module that implements the Python multiprocessing library and allows to execute remote processes in serverless functions in a transparent way. The project is open source and the source code is publicly available on Git-Hub.<sup>2</sup>

The open questions presented in this article are: Can we run legacy single-machine parallel applications in the Cloud and scale them transparently using serverless resources? Can we program the Cloud as an infinite multi-core machine?

## 2. Related work

**OS-level transparency.** Hiding the complexity of distributed systems is a recurring topic in the systems field. Recent industrial trends on Disaggregated Data Centers (DDC) [5] advocate for a distributed OS transparently leveraging disaggregated hardware resources like processing, memory or storage.

<sup>2</sup> <https://github.com/lithops-cloud/lithops>.

For example, LegoOS [13] is a disaggregated OS that implements a subset of the Linux system call interface so that existing unmodified Linux applications can run on top of it. LegoOS shows how two unmodified applications can be run in a distributed way: Phoenix (a single-node multi-threaded implementation of MapReduce) and TensorFlow. LegoOS is however not demonstrating scalability or complex scenarios involving mutable memory. Even worse, implementing the entire Linux APIs over disaggregated resources is a daunting engineering task.

Another approach in access transparency at the operating system layer is GiantVM [14]. GiantVM uses virtualization to run an unmodified guest OS over a distributed cluster. In contrast to the traditional many-to-one virtualization paradigm (running multiple OS in one machine), GiantVM implements one-to-many virtualization (running a single OS in many machines). GiantVM uses *Infrastructure as a Service* (IaaS) to run the distributed OS over a cluster in a Cloud setting. Approaches based on current Cloud technologies are more feasible at the moment, since DDCs are not yet available to the general public.

In [6], the authors propose to augment operating systems for disaggregation, by exposing explicitly the disaggregated resources to applications and thus opening efficient and optimized co-designs between applications and the remote resources.

**Application-level transparency.** Instead of at the operating system level, transparency can also be achieved easily at the application level. If the interface that is used by the application to access local resources is replaced or wrapped by another implementation that accesses disaggregated resources instead, we could then transparently run unmodified local code in a distributed fashion.

Early efforts in the Serverless community propose *FaaSification* as an automated process to move existing code to Serverless Functions [15]. Although this only implied simple functional code, and not entire applications.

Fiber [16] is a library that implements Python's multiprocessing API to run remote processes in a distributed Kubernetes cluster. In their article, they execute several stateful AI applications that are programmed for local parallel execution using Python's multiprocessing library. By replacing multiprocessing with Fiber, those unmodified applications can transparently scale and exploit parallelism on a distributed Kubernetes cluster.

Although our work is closely related to Fiber, they compare Fiber with other distributed computing frameworks, while we wanted to focus on studying access transparency and to find the differences between local and distributed execution exploiting the high scalability of FaaS. Moreover, Fiber does not implement all Python multiprocessing abstractions (for example, it misses the `Lock`).

Another example of transparency at the application level is Crucial [17]. Crucial is a library that implements Java's threading interface which allows threads to be transparently executed as serverless functions. Using serverless, Crucial highly improves scalability thanks to the inherent elasticity of FaaS. Crucial also provides stateful abstractions based on distributed shared objects that reside in a disaggregated in-memory layer (Infinispan). Nevertheless, Crucial is not offering transparency to Java applications, since it offers an explicit programming model for shared remote objects.

Kappa [18] is a serverless computing framework that focuses on providing fault tolerance for stateful serverless applications by means of checkpointing mechanisms and continuation functions. Contrary to our work, they do not emphasize access transparency in their contributions. Yet, they claim that their framework requires minimal modifications to the original code, because Kappa's static code analysis system is able to

create checkpoints automatically. However, a special API is required to invoke tasks (`spawn` and `spawn_map`) and to pass messages between functions. We believe it would be interesting to study fault-tolerant access transparency with Kappa framework in future work.

We differ from related work in that we are the first work that evaluates full transparency over serverless disaggregated resources intercepting multi-processing parallel libraries.

### 3. Enabling transparency for Python multiprocessing

This section describes how we achieved access transparency to serverless and disaggregated resources through the re-implementation of the Python multiprocessing library. Fig. 2 depicts a general overview of the architecture.

We have leveraged the Lithops framework [12] to execute parallel local applications over disaggregated serverless functions. Lithops enables the execution of local serial code to be run over massively parallel serverless functions. Lithops acts as an abstraction layer that simplifies the exploitation of the main FaaS services present in public clouds for highly parallel tasks. One of Lithops design principles is to ensure portability between clouds. The same application can be seamlessly ported from one cloud provider to another, which prevents vendor lock-in.

We have extended Lithops with a multiprocessing module which implements in its entirety the original Python multiprocessing interface. Computation abstractions (like `Process` and `Pool`) use Lithops `FunctionExecutor` API. Inter-process Communication (IPC) and synchronization abstractions (like `Lock` and `Queue`) are implemented using Redis key-value in-memory database.

#### 3.1. Disaggregated compute resources

##### 3.1.1. Lithops workflow

Lithops follows a main/worker architecture where a local process acts as the orchestrator and coordinator of the workers that are deployed and executed in the FaaS backend as serverless functions. A diagram of the general operation of Lithops is shown in Fig. 3.

First, the user interacts with Lithops multiprocessing API, which is a wrapper around Lithops framework API (1). Lithops automatically detects, serializes and uploads to storage the processes' dependencies, process function code and input arguments (2). Next, the Lithops orchestrator invokes the corresponding number of serverless functions against the FaaS backend (3). For example, every `Process` corresponds to a single function. The Lithops worker (4) is a generic serverless function that handles the execution of Lithops job tasks. It downloads the previously uploaded code, data and dependencies from storage, deserializes them and it executes the user's function in a wrapper that handles errors. When the task is completed, the result is uploaded back to storage. The Lithops orchestrator synchronizes completed tasks by pulling the contents of the storage (5) — a task has finished when the result key is listed. Finally, the results are downloaded and returned to the parent application.

##### 3.1.1.2. Serverless job queue

Forking many local processes for very fine-granular and short-running tasks is slow, and is much worse in Lithops, as the overhead of invoking many functions can be prohibitively expensive. For this reason, Python multiprocessing implements the `Pool` abstraction. A `Pool` represents a fixed-size group of processes that are forked at the time of the pool instantiation. It has methods like `starmap()`, `apply_async()` or `map_async()` which allows to offload tasks to the worker processes at a higher

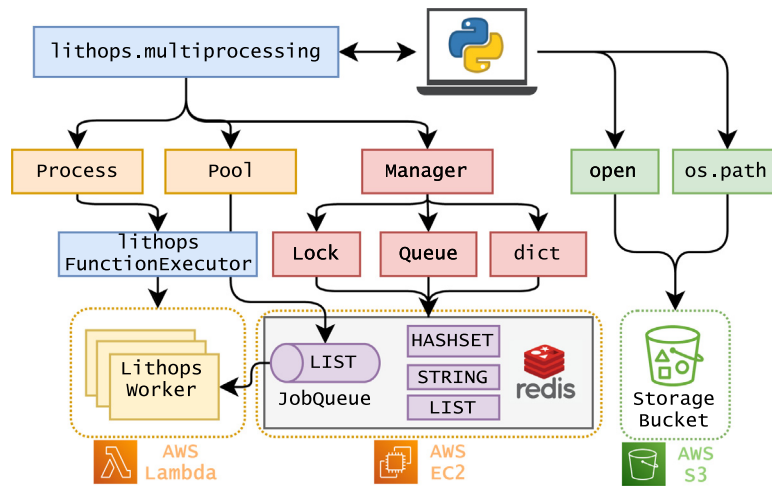


Fig. 2. Architecture diagram.

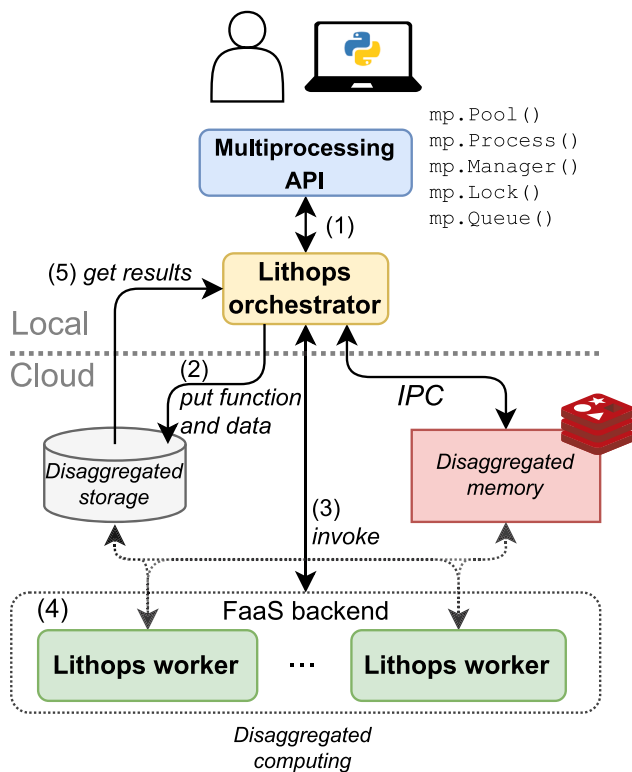


Fig. 3. Lithops operation workflow.

level, instead of manually creating and managing Process instances. In a process pool, each operation (*map*, *apply\_async*...) creates one or more jobs, which are enqueued to a job queue. The worker processes get and execute jobs from the queue. This avoids the need to create new processes for each task, which considerably reduces the fork overhead. In addition, it is useful to initialize global worker-scope variables, as they only need to be initialized once at worker process creation.

We have implemented the job queue pattern for the Lithops multiprocessing Pool. In `Lithops.multiprocessing.Pool`, workers are long-lived functions that are invoked when the Pool object is created. Operations on the pool (`map()`, `apply()`) generate Lithops tasks, but instead of invoking new functions to execute those tasks, they are queued in a Redis list. The worker

functions pick up and execute tasks from the queue as they are generated. Once the Pool is closed (`terminate()`), a message is sent to the workers to terminate their execution.

The main advantage of this implementation is that the overhead of submitting a set of tasks to a Redis list is much lower than invoking a function for every task. With Redis, we can submit all tasks at once with a single LPUSH command, while invoking functions is sequential and the overhead depends on the API and architecture of each FaaS service. Also, reusing functions to execute multiple tasks avoids stragglers caused by cold invocations. However, the main drawback is the function execution time limit. Although the time limit has been increasing over the years (for example, AWS Lambda now supports invocations of up to 15 min [19]), this limit prevents running longer executions.

### 3.2. Disaggregated memory resources

We have chosen Redis for its simplicity of deployment, in-memory storage and high performance. Redis differs from other traditional key-value databases in that the value has a type, such as LIST, STRING, HASHSET, etc. The different operations available on these data types facilitate the implementation of some of the communication and synchronization abstractions present in Python’s multiprocessing library. A single node Redis deployment guarantees consistency and the correct order of read and writes, since Redis is single-threaded, and data is backed in disk for restart recovery, but node failure tolerance is not provided.

Serverless functions are not addressable and lack the ability to open a connection to one from another. To circumvent this limitation, many works have explored using TCP connections through an addressable external proxy residing in a VM [9,20], or using NAT hole-punch techniques to break FaaS networking and enable a P2P-like connection between functions [21]. However, all these approaches still need an external rendezvous server, but most importantly, they do not support collective communications. Many multiprocessing abstractions require synchronized and atomic access, such as queuing an item to a Queue or signaling the release of a Lock, and Redis single-threaded implementation meets this requirements in a safe but fast manner.

In Python’s multiprocessing, processes need a reference to the objects that represent a shared state resource (such as a Queue or Manager instance). The parent process creates all resources and then it passes a reference to the child processes when they are forked. With Lithops, objects and references passed to functions have to be serializable. To maintain the same behavior, we have followed a pattern in which each resource object (Queue,

Pipe...) acts as a proxy to the key–value pair in Redis, which is where the state resides. Each object is uniquely identified and corresponds to a specific Redis key–value pair. Each proxy resource implements reference counting for garbage collection. The counter is consistently stored in Redis, and the resource is deleted from Redis when references reach zero. In addition, each resource incorporates a key expiration time of an hour by default. It is used as a backup to in case there is an error in the code and the program terminates abruptly, since the reference counting mechanism might not delete the resource in a graceful manner.

The different abstractions available and a brief description of how they have been implemented are listed below:

**Message passing:** *Pipes* are used for duplex communication between a pair of processes. A new *Pipe* returns a tuple of two *Connection* objects, each corresponding to a Redis *LIST*. Data can be written to one end of the *Pipe* using *Connection.send()* and it can be read using *Connection.recv()* on the other end. The *send()* method executes an *LPUSH* command to put data in the tail of the list, and the *recv()* method executes an *BLPOP* command to get data from the head list. This way, the list is treated as a FIFO queue. The *BLPOP* gets and removes the head item of the list, or blocks until there is an element available. *Queues* are implemented the same way as *Pipes*, the difference being that more than two processes can put or remove items from a queue. Redis maintains the order of puts and gets consistent.

**Shared state:** *Array* and *Value* are used to share memory in Python multiprocessing. Only basic C type values can be put into the array or value. They are implemented using the *LIST* type. Processes can read and write to specific indexes or slices of the array. A *Value* is an *Array* of size 1. We have opted for using the *LIST* type instead of *STRING* because *STRING* values are limited to 512 MB in size, while in lists, each element of the list can will be at most `sizeof(long double)` in size, and lists can hold up to  $2^{32} - 1$  elements.

*Managers* allows to create Python resources in a separate process, and they are accessed via sockets. *Managers* are used to share a basic Python data type with multiple processes, such as a *dict* or *list*. The implementation of those types is trivial using Redis, since it provides *HASHSET* and *LIST* types natively. A *Manager* also permits the creation of user-defined classes, which reside instantiated in the *Manager* and other processes use *Remote Method Invocation* to access it. To provide a similar behavior using Redis, we have made each process have a local instance of the *Manager* class, but the state of the user-defined class instance (i.e. its attributes) is remotely stored in Redis as simple key–value pairs. A *Lock* ensures that attributes are accessed by only one process at a time.

**Synchronization:** *Semaphores* and *locks* are implemented using the *LIST* type. When a *Semaphore* is created, *N* tokens are added to the list, being *N* the initial value of the semaphore. Every *acquire()* of the semaphore will execute a *BLPOP* command, which removes a token from the list. The *release()* method puts the token back to the list with a *LPUSH* command. If there are no tokens when a *acquire()* is called (meaning that *N* processes are currently in the critical section), the *BLPOP* command will block until other process performs a *release()* and puts back a token into the list. Note that in this implementation, the value of the semaphore will always be greater than zero. A *Lock* is a generalization of a semaphore where *N* is 1.

To implement *Conditions*, multiple *notification lists* are used to notify blocked process awaiting for the condition event. When a process reaches the condition and hangs on the *wait()* method, the process registers a new list to the *notification list set* and blocks to it with a *BLPOP* command. The process that satisfies the condition will add an element to each list of the *notification list set*, so that all waiting processes are unblocked and the execution resumes. *Barriers* and *Events* are specific cases of *Condition*.

### 3.3. Disaggregated storage resources

Lithops multiprocessing also implements a replica of Python's built-in *open* function and the *os.path* module which allows to transparently read and write files and directories stored on disaggregated storage services (like *S3*) as if it were a local file system. This is especially useful for *FaaS* since the volume that is mounted in the function container is volatile and the data stored there is lost when the execution finishes. In this way, we offer serverless processes a transparent way to save or recover their state. However, it should be noted that, since we are working on immutable data, it is not possible to modify or expand a file as would be done in a traditional network file system without having to rewrite the entire file, which can be problematic for large files. However, as seen in Section 5.4, disaggregated storage services provide much higher parallel read and write throughput than traditional disks used in monolithic machines. Applications that require reading lots of data in parallel (for example, video encoding) can benefit from disaggregated storage to achieve lower execution times.

## 4. Evaluation settings

This section aims to describe the configuration with which the experiments described below have been carried out.

Unless otherwise stated, the experiments have been run with the following settings: Lithops orchestrator runs on a m5.2xlarge EC2 host with Ubuntu 20.04, Lambdas use a containerized Python 3.8 runtime with 1769 MB of RAM<sup>3</sup> as serverless function and Redis 6.2 instance runs on the host machine with Docker.

The host machine and the AWS Lambdas are in the same VPC private subnet, region and availability zone (us-east-1 A), so traffic does not go through a NAT gateway nor the public internet. In case of using *S3* as storage backend (will be stated), the *S3* bucket is located in the same region and access to *S3* is done via a private endpoint. All Lambda functions have been executed using warm containers.

All local monolithic executions have been carried out using on-demand AWS EC2 instances with different number of vCPUs. In particular, we have used the following instances: c5.4xlarge with 16 vCPUs, c5.9xlarge with 32 vCPUs, c5.18xlarge with 64 vCPUs and c5.24xlarge with 96 vCPUs. All of these c5 EC2 instances are running Ubuntu 20.04 and are located in the us-east-1 region. Note that the shared memory abstractions for local executions use local memory, while remote serverless processes use Redis.

Since we want to study access transparency, the code used for both configurations is exactly the same, except that we change the import statement so that serverless executions use remote stateful abstractions (from `import multiprocessing` to `import lithops.multiprocessing`). In other words, all local executions with a single virtual machine make use of the standard Python multiprocessing library and local resources (processes, shared memory), while serverless executions use the Lithops multiprocessing library and remote resources instead (serverless functions, Redis).

We have discarded to perform executions in an on-premise environment for two main reasons. First, we cannot emulate a Cloud environment using physical resources, since we do not certainly know the hardware used in EC2 and because of Cloud multi-tenant resource-sharing. Second, AWS Lambda runs on EC2 type instances, so comparing local on-premises executions with Lambda would be unfair.

The source code for the different validations is open source and publicly available on GitHub.<sup>4</sup>

<sup>3</sup> According to AWS documentation [19] a runtime of 1769 MB of memory is assigned a whole vCPU, being a vCPU a thread of a CPU with Hyper-Threading [22].

<sup>4</sup> <https://github.com/aitorarjona/lithops-transparency-validation>.

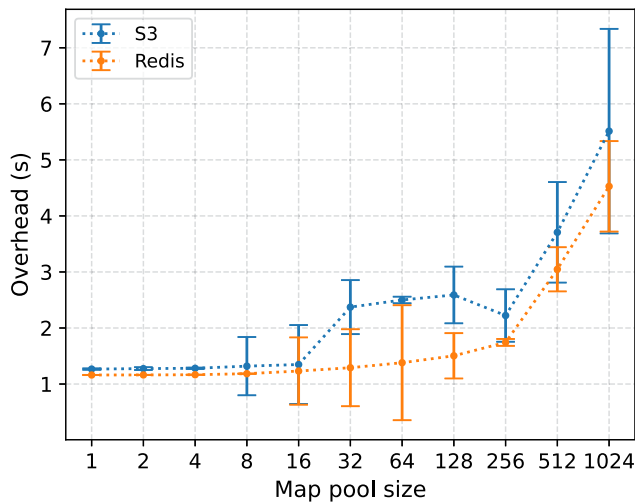


Fig. 4. Processes pool invocation results.

## 5. Micro-benchmark evaluation

The objective of this section is to perform micro-benchmarks that allow us to identify potential overheads and where are they generated, which will help us understand the validation results of the real applications.

### 5.1. Fork-join overhead

One of the main overheads when doing parallel computing is the cost of creating a new thread or process. The purpose of this experiment is to measure the overheads generated when invoking multiple parallel serverless processes and analyze how they scale when the number of functions is increased. Lithops allows the usage of different storage backends and monitoring systems for serverless functions. In this experiment, we will compare the performance between using S3 or Redis as storage backend and task monitoring.

The experiment consists of performing a multiprocessing Pool of several sleep functions, each one would sleep for 5 s.

Fig. 4 represents the overall overheads for different parallel executions. The overhead time is calculated by subtracting the sleep time from the total execution time. For low parallelism, both Redis and S3 synchronization provide an overhead of  $\approx 1$  s. From 32 processes and up, Redis can provide lower overheads than S3: At 1024 functions, Redis synchronization provides an overhead of  $\approx 1$  s lower than S3.

We also have measured where these overheads are generated by Lithops and AWS Lambda invocation.

Fig. 5 represents a breakdown of a map job of 5 sleep seconds for 1024 parallel functions using Redis as storage backend. The execution shown in the left chart used cold containers, while the one on the right used warm containers. We can see the difference in the function start-up overhead. When using cold containers, the overhead is higher because the provider has to allocate resources to run the functions, while when using warm containers, the resources are already allocated and the container is already up and running, so the overhead is much lower. Specifically, warm invocations typically have an overhead of around 200 ms, while cold invocations can have a more variable and higher overhead. In this case, the overhead of cold invocations exceeds the second and the average overhead obtained is 1.7 s. We can also see that the start of execution is not instantaneous but linear. This is because asynchronous invocation using Python threads is performed sequentially. This implies that the greater the number of

Table 1

Decomposition of generated overheads. The values indicate the average across all functions of a map job.

Phase	Invocation type	
	Cold	Warm
<i>Serialize data and function</i>	0.004 s	0.004 s
<i>Upload dependencies</i>	0.002 s	0.001 s
<i>Invoke</i>	1.719 s	0.258 s
<i>Function setup</i>	0.052 s	0.046 s
<i>Join</i>	0.628 s	0.630 s
<b>Total</b>	2.407 s	0.939 s

Table 2

Decomposition of generated overheads. The values indicate the average across all functions of a map job.

Payload size	Remote (Redis)	Local
1 KB	0.6 ms	0.0463 ms
1 MB	23.4 ms	2.56 ms
100 MB	1.12 s	0.288 s

functions, the greater the invocation overhead. It also implies that full parallelism is not achieved immediately. Applications that require exact parallel execution should use some synchronization mechanism (e.g. a barrier).

Table 1 shows the decomposition of the overhead introduced by Lithops and by AWS Lambda. The values indicate the average times of all the functions of the same map job. We have differentiated two executions, using warm and cold containers. “*Serialize data and function*” and “*Upload dependencies*” indicates the time spent to serialize and upload the function data and its dependencies, respectively. These values are constant since both use the same data and storage backend. The “*Invoke*” row indicates the time elapsed since the function is invoked until the function begins its execution, and “*Function setup*” indicates Lithops worker wrapper setup time. We can see that it is longer for executions using cold containers for the reasons exposed above. The “*Join*” time indicates the time elapsed since the function ends until it is detected by the Lithops orchestrator. Finally, the sum of all overheads is shown in the “*Total*” row.

The overhead time determines the minimum granularity of process tasks. Short-running tasks with run time lower than the overhead will not benefit from distributed serverless execution. Moreover, the closer the granularity is to the overhead time, the more noticeable it will be with respect to the total application execution time.

### 5.2. Network latency and throughput

The main limitation and bottleneck for access transparency will be the access latency to remote shared stateful resources. With these experiments, we want to determine the limits of shared memory using Redis and the latency and bandwidth of the network between remote processes.

First, we will determine and compare the latency for local (VM) and remote (serverless functions) communication. We send a variable size payload through a multiprocessing Pipe and measure the round-trip time. Local Pipes are UNIX pipes while Lithops Multiprocessing Pipe uses Redis Lists. Results on latency are arranged on Table 2. We see that the latency of remote communication using Redis is an order of magnitude higher than local communication, so the performance is not comparable. However, we note that for small payloads (less than 1 KB), the latency is below a millisecond, which makes the overhead of synchronization operations (like Locks or Semaphores) that do not require data passing of minor relevance.

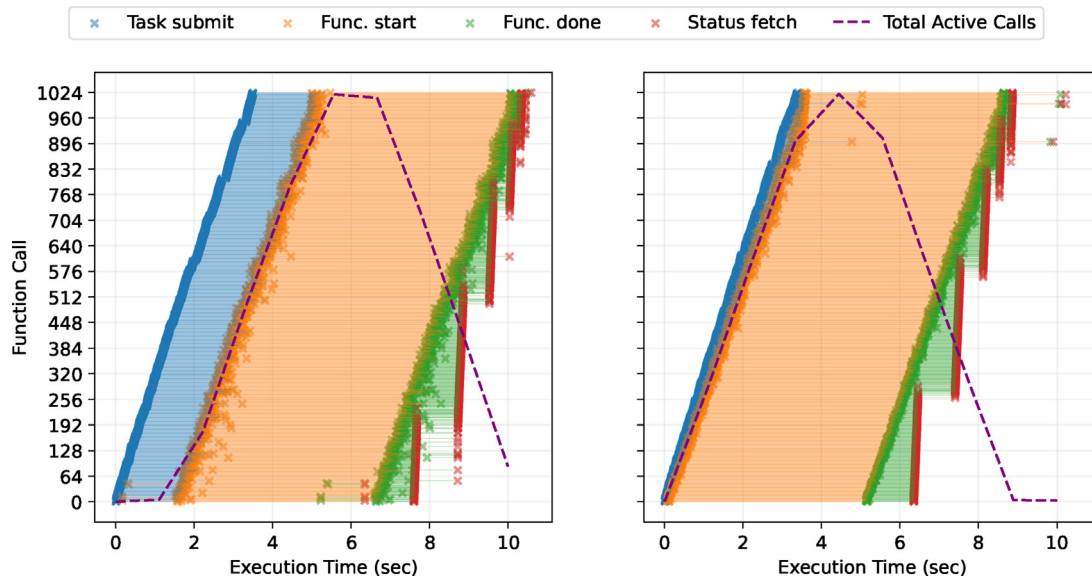


Fig. 5. Breakdown of the function's stages for the `sleep(5) map` job. (a) Cold invocation (b) Warm invocation.

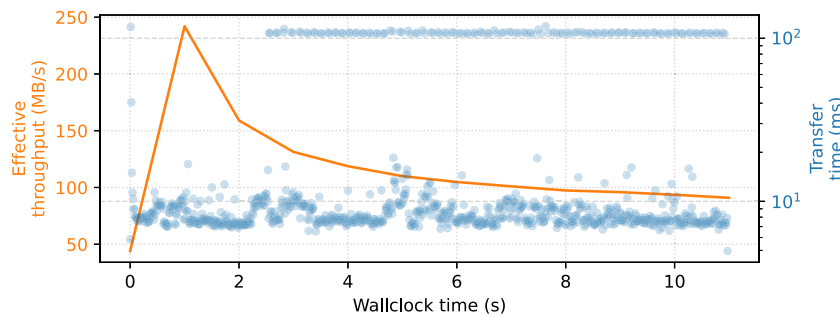


Fig. 6. Throughput results.

Second, we want to measure the maximum throughput of a Lithops multiprocessing Pipe backed by Redis. We send 1000 messages with a size of 1 MB (for a total of 1 GB) through a Pipe to communicate two remote processes running in serverless functions. We can see in Fig. 6 that the time elapsed for sending a message is stable at 15 ms, although we can observe some outliers, which could be caused by shared network usage interferences. The total transmission takes 10.5 s, so the effective throughput rate is  $\approx 90$  MB/s. From this result we can determine the transmission data time between remote processes, which will indicate whether it is worth using disaggregated resources for certain workloads.

### 5.3. Computational performance

The goal of this experiment is to measure computational performance in an embarrassingly parallel example and to compare the execution time and scalability between a large VM and serverless functions with Lithops.

We used the classic example of the calculation of Pi with the Monte Carlo method to carry out the experiment. In particular, this test is based on sampling 3,200,000,000 random points and calculate the number of points that are within the unit circle to extract an approximation of the Pi number. The amount of points to sample is distributed between all the processes, so execution time should decrease when increasing the number of processes.

The results of Fig. 7 show that the scalability capacity that can be obtained with Lithops using FaaS goes much further than

what a single machine could achieve, despite the executed code is exactly the same. The baseline execution time is 1254.8 s for a single process. For 16 processes, we observe that the performance of the disaggregated system is between 20% and 25% superior compared to the monolithic system. This is caused because, in a VM, the processes share physical CPU cores between them thanks to the use of Hyper-Threading technology, but despite this, the number of floating-point operations is still limited by the physical CPU cores used. On the contrary, functions do not present this limitation, because the physical nodes used by AWS Lambda have Hyper-Threading disabled [23]. However, if we adhere to the documentation, for 1769 MB of memory, the function is allocated one vCPU [19], i.e. one CPU thread [22]. If we refer to the official documentation, we can conclude that, in equivalence of number of vCPUs, functions have a computational advantage over VMs. We can also observe that, for 96 processes, i.e. the VM ceiling, both disaggregated and monolithic systems converge and present approximately the same performance. This is caused because Lithops overheads (see Section 5.1) are masked by Hyper-Threading inefficiencies present in VMs.

### 5.4. Disk performance

The objective of this experiment is to measure the disk read and write capacity and scalability for Lithops' processes running on FaaS, transparently emulating the disk of a VM.

The experiment has two phases: In the first phase, a batch of processes will write a 1 GB file to disk. In the second phase,

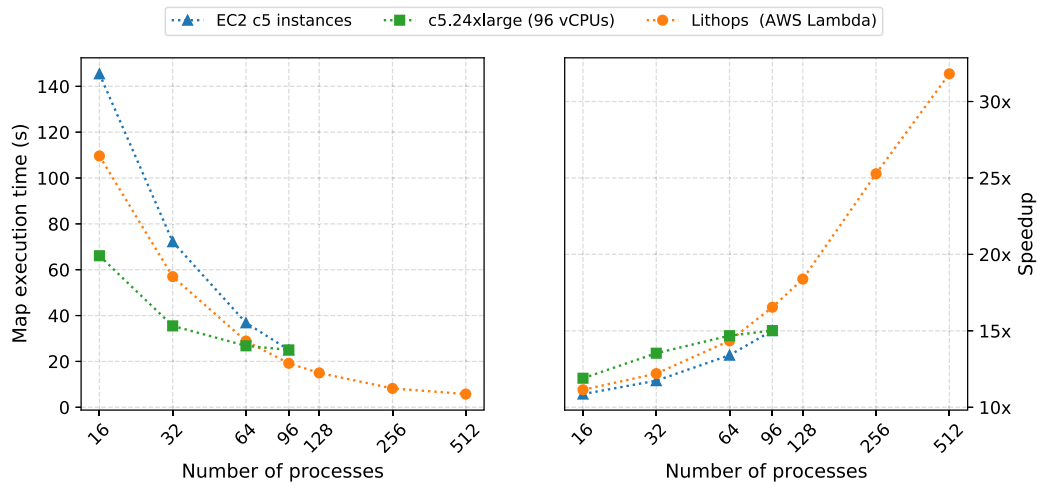


Fig. 7. Pi Monte Carlo results.

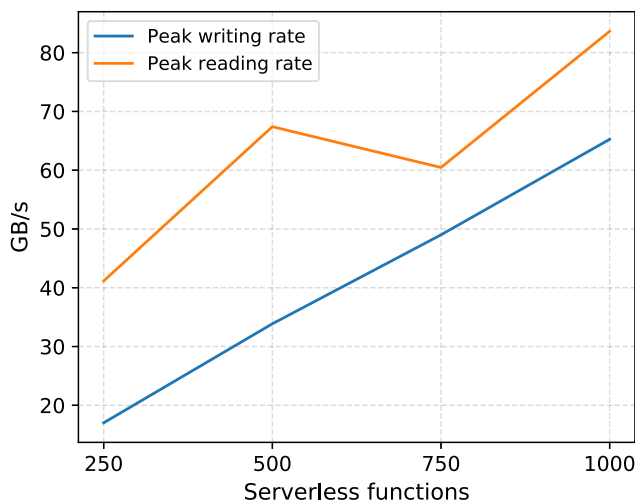


Fig. 8. Disk reading and writing rates.

another batch of processes read from disk the data written in the first phase. The experiment is performed for different numbers of processes to study scalability, where we measure aggregate read and write rates. The storage backend used is S3.

The results, depicted in Fig. 8, show high scalability and aggregate bandwidth, with peaks of 80 GB/s for reads and 65 GB/s for writes. Compared to General Purpose SSD (gp2) EBS volume, which its maximum throughput is 250 MiB/s<sup>5</sup> [24], the aggregate write/read throughput of disaggregated storage is considerably higher compared to that of a single volume mounted on a VM. This may be useful for applications that require reading lots of data from storage in parallel since with FaaS we can achieve higher throughput and thus a lower execution time.

### 5.5. Shared memory performance

We also want to validate the performance loss when using remote shared memory. For this experiment, we have implemented a parallel sorting algorithm using shared memory and compared local and serverless execution. The sorting algorithm consists of splitting in chunks an array, performing quick sort on each chunk in parallel, and then recursively merge pairs of

Table 3

Execution time of the different parallel quick sort implementations with different array size.

Array size	5 M		10 M	
	Local	Lithops	Local	Lithops
Shared array	23.66 s	–	79.68 s	–
Shared array with local copy	15.693 s	356.60 s	47.11 s	–
Message passing	14.27 s	17.30 s	45.16 s	45.63 s

sorted chunks following a tree pattern. This algorithm makes heavy usage of the array as it is iterated over multiple times and the list items are constantly changing positions. We have followed three strategies to implement this algorithm. The first uses multiprocessing’s shared Array to store the array and the operations are performed in-place, directly accessing the array indexes. Although the array is stored in shared memory, each process only accesses its corresponding chunk, so there is no need to provide mutual exclusion and critical sections. The second strategy also uses a shared multiprocessing Array, but each process copies its chunk to a local variable, performs the operations on this local slice and then copies back the slice to the shared memory array. The third strategy does not use a shared array, but Pipes and message passing instead. The parent process sends the chunks to each worker process using pipes, and the workers perform the tree merge phase also passing their chunks using pipes between them.

We have implemented these three strategies to show that, although all three are correct for local execution, the way memory is accessed has a huge impact on the performance when using disaggregated memory.

Table 3 shows the execution times of the three alternative implementations using different array sizes (5 M and 10 M) on both local and serverless using 64 processes. For the local execution, we have used a c5 EC2 instance with 64 vCPUs. We can see that Lithops was not able to execute the algorithm using the in-place shared array implementation. This is because each access to a list index is equivalent to a Redis command request. This causes a prohibitive overhead that prevents obtaining competent results using this shared memory approach. The local copy implementation is presented as a low-effort improvement over the in-place shared array implementation. Still, Lithops struggles to perform due to the high data copy overhead. Finally, the message-passing approach using Pipes is presented as the proper way to implement this algorithm using disaggregated memory. Since shared memory is no longer used, Lithops is now capable of

<sup>5</sup> for volumes bigger than 334 GB regardless of burst credits.

**Table 4**  
Summary of algorithm type and stateful abstractions used for every application.

Application	Algorithm type	Shared state	Message passing	Parallel map
Evolution strategies	Iterative pool map	✓	✗	✗
Pandaral-lel	Scatter-Gather	✗	✓	✗
Gridsearch	Broadcast-Gather	✗	✗	✓
PPO	Main-Worker	✗	✓	✗

providing competent performance compared to local execution. For the 10 M array size execution, we can see that both local and serverless executions result in the same execution time. We also can see an improvement in the local execution compared to the shared memory alternative. This tells us that, even if we have fast-access local shared memory, it is sometimes not the best alternative even in local executions.

### 5.6. Micro-benchmarks conclusion

As we have seen in the different experiments carried out, the overheads, mainly generated by network latency, are very considerable. We are obtaining overheads of several magnitudes higher (microseconds locally compared to milliseconds remotely). However, the overall performance is not severely affected, because the Hyper-Threading penalty masks the overheads originated by the network. Using FaaS, such as AWS Lambda, gives us great instantaneous scalability, although the additional overheads are more significant. For example, the lack of direct communication between functions makes it necessary to use indirect communication via a disaggregated memory component. In this regard, we are evaluating transparency in a non-optimal scenario — a better alternative would be to use multiple VMs, where we would benefit from direct communication or where we could avoid the function invocation overhead (with already collocated remote processes). On the other hand, we would lose the ability to scale up quickly and dynamically. Precisely, these attributes are of vital importance when it comes to transparently running Python scripts that use multiprocessing. Until the program is in runtime, we generally cannot know in advance the number of processes or which shared state communication abstractions we are going to use, or whether the number of processes varies throughout program execution. If we were to use a cluster of virtual machines, the total capacity would be fixed, so there would be times when over-provisioning (having more resources deployed than we really need) or under-provisioning (having fewer resources that are throttled by a greater load) would occur.

## 6. Applications

In this section, we evaluate the behavior of Lithops in four different real use cases in order to test access transparency and to measure performance. The scenarios used are: the implementation of the OpenAI's Proximal Policy Optimization (PPO) algorithm in its Baselines repository, the POET modifications in Evolution Strategies made by the Uber research team, parallel Pandas dataframe transformations using Pandaral-lel, and a hyperparameter tuning using Scikit-learn's Gridsearch running over Joblib. To adapt these applications to serverless, we only had to replace the `multiprocessing` import with `Lithops.multiprocessing`. Since Lithops fully implements the multiprocessing interface, the rest of the code did not need any further modification.

Table 4 describes the applications and the type of algorithm used in them. We also specify what kind of stateful abstractions

are used. In each section, we go into detail about how each algorithm handles the shared state and message passing.

We have taken the PPO and Evolution Strategies applications from the Fiber [16] validation since they are real and complex applications that use the multiprocessing library. However, we believe that the comparison with fiber would not be fair, due to (i) container creation times in Kubernetes and AWS Lambda are not comparable and (ii) Fiber uses direct inter-process communication while Lithops uses disaggregated memory. Moreover, they do not indicate enough parameters to replicate their experiments.

Note that the use of these experiments helps us to have complex scenarios in which to check if access transparency can be achieved, in no case it is intended to study the results of the experiments in the fields of artificial intelligence, machine learning or data analysis. Note that as each application is different, so is its scalability, therefore the number of processes used in each application varies.

### 6.1. Evolution strategies

In this experiment we have used the Paired Open-Ended Trailblazer (POET) [25] implementation, which is a large Python application with about 4000 LOC (lines of code) using different *multiprocessing* abstractions like `Pool` or a shared dictionary from a `Manager` (`Manager.dict()`). This algorithm is part of the Evolution Strategies category, in which, evolutions of an initial population are carried out iteratively, and those evolutions are executed in parallel. The objective of this test is to analyze and compare the performance and scalability of Lithops in an iterative algorithm that maintains and uses a shared state between processes.

POET uses a shared noise table which is used to generate randomness in the evolution process. This noise table is originally implemented using shared memory. However, it is initialized when the module is loaded, so it is not using Lithops multiprocessing implementation for shared memory. Instead, since this table is read-only, each function can initialize its noise table independently of the shared memory. The algorithm also uses a shared table of parameters that are modified in each iteration. This shared data structure is implemented as a shared multiprocessing `Manager.dict()` dictionary. Therefore, there is a certain transmission of data that could imply a significant overhead.

The *multiprocessing* abstractions used in POET are: one `Context` set to `spawn` mode, one `Pool` for tasks executions, one `Manager` with two `Dict` that contain the shared stated used by all the worker processes from the `Pool`.

Each iteration of the algorithm performs a `Pool.map()` operation. As the task granularity is small (about 3 s), to try to mitigate overheads, we used the optimization of the `Pool` with job queue explained before. To carry out the measurements we have executed 5 iterations with 512 batches per chunk and a batch size of 5. All local executions have been run on a `c5.24xlarge` EC2 instance.

The results in Fig. 9 show that, despite the data transmission and invocation overheads, Lithops maintains constant scalability similar to the scalability of the VM. The maximum speedup of the VM is about 40x, while Lithops is capable of reaching a speedup of around 53x, improving the best result of the VM.

### 6.2. Pandaral-lel

Pandas is a Python package, with more than half-million lines of code, that provides data structures to work with relational data. Pandaral-lel [26] is a Python module with about 1700 LOC that extends the Pandas functionality adding parallel `DataFrames` and `Series` transformations functions like `apply`, `map` and `applymap`.

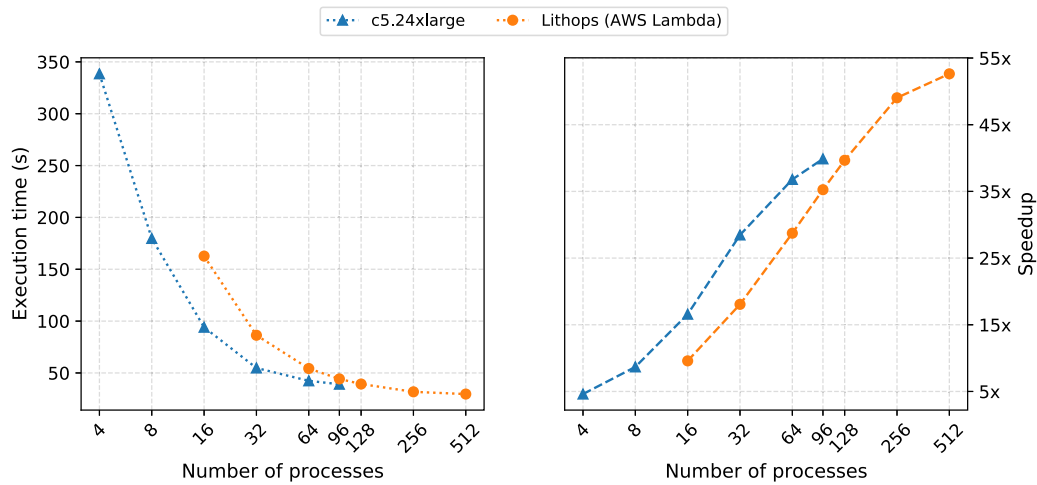


Fig. 9. Evolution Strategies execution results.

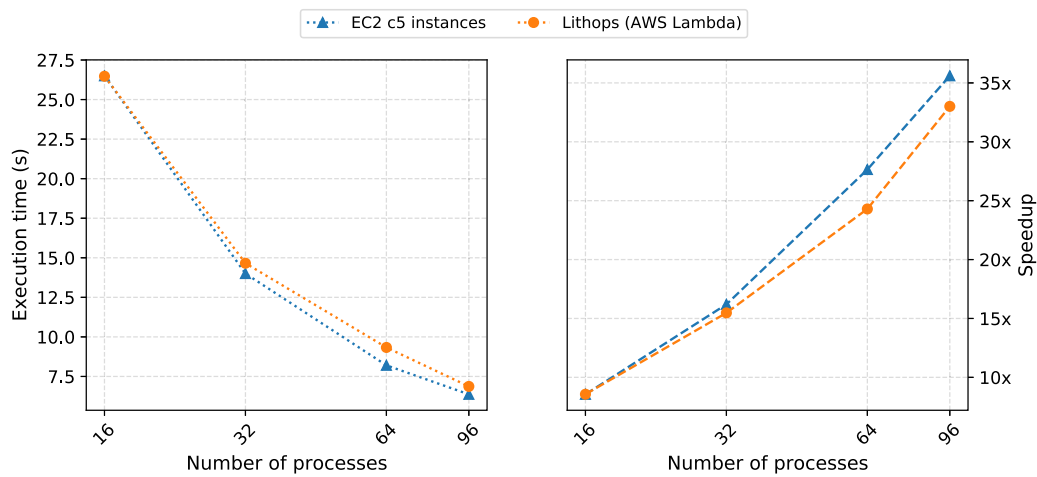


Fig. 10. Pandaral-1el execution results.

To do so, Pandaral-1el relies on the multiprocessing API, so replacing it by the `Lithops.multiprocessing` API allows us to execute it in a distributed environment. This experiment aims to measure Lithops' behavior in an embarrassingly parallel task with relatively large data transmissions and analyze how it handles that overhead.

For this experiment we used the Sentiment140 [27] dataset loaded in a Pandas DataFrame. We used Pandaral-1el `apply()` function on that DataFrame to perform some sentiment analysis using the `textblob` Python module. Pandaral-1el first partitions the dataframe according to the number of available workers. Then, it serializes the content of each one and passes it to each function by parameter. Lithops detects the values passed by parameter of each function and transfers them to the storage. The functions are then invoked using `Pool.map()`. Each function downloads its chunk of dataframe, applies the transformation and, finally, the resulting dataframe is returned to the parent process. During this process Pandaral-1el uses the following multiprocessing artifacts: one `Context` set to `fork` mode, one `Manager`, one `Pool` for tasks execution and one `Queue` to synchronize the main process with the workers.

The results in Fig. 10 show that Lithops obtains a 7% lower performance compared to the best result from the VM. It can also be seen that Lithops is capable of maintaining correct scalability up to 96 vCPUs, which can be deployed immediately without previous allocation.

There are two main reasons for these results. The first reason is data transmission and partition. Notice that the dataset used is above 200 MB on disk (and about 600 MB when it is loaded into memory) that has to be transmitted from the Lithops orchestrator entirely to the storage backend. The second reason is task granularity. In the results can be observed that the differences between VMs and Lithops grow as the number of vCPUs increases. The reason behind this is, as seen in Section 5.1, fork-join overhead grows as the number of parallel processes augments. When the granularity of the task is too small (note that with 96 vCPUs the total execution time is below 7 s), the Lithops overheads become a significant percentage of the total execution time. For this reason, in 16 or 32 vCPUs the difference between Lithops and VMs is much smaller than when using 64 or 96 vCPUs.

### 6.3. Scikit-learn hyperparameter tuning

Scikit-learn is a Python module for machine learning built on top of SciPy. It is widely used and considered a standard to build machine learning models to solve classification, regression and clustering problems in the Python ecosystem. Here we want to test Lithops behavior on an embarrassingly parallel scenario with low data transmission.

Some Scikit-learn utilities can parallelize their execution via the `joblib` [28] library. Joblib provides a lightweight pipelining in Python, in particular, an API to do easy simple parallel computing.

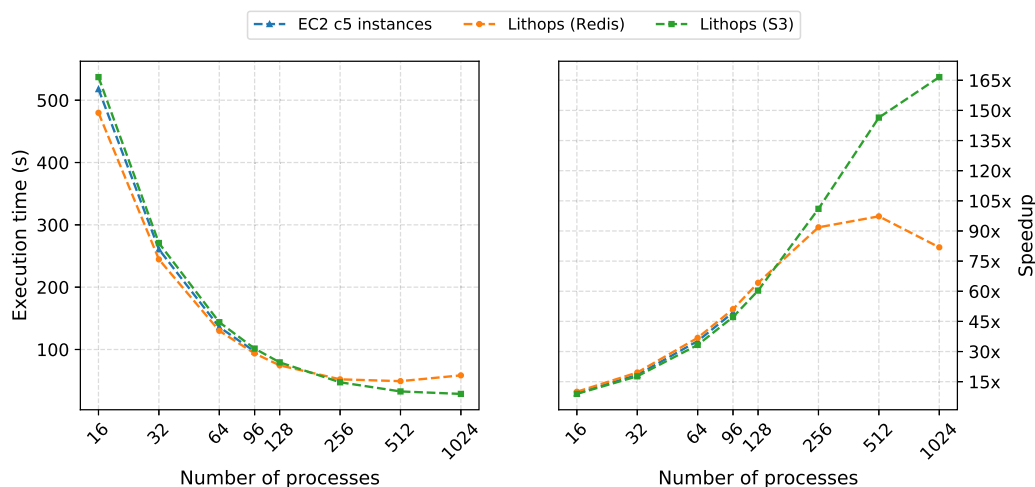


Fig. 11. Hyperparameter tuning execution results.

Joblib allows using different parallel backends such as *loky* or *multiprocessing*, but it also allows you to use and develop your own parallel backend. We have created a Lithops backend for joblib using the original *multiprocessing* backend as a template. Changes with respect to the original *multiprocessing* parallel backend are minimal. To use the Lithops backend, the user has to first register this backend into Joblib, and then select it to run the parallel jobs. After that, Scikit-learn transparently handles all job lifecycle while the original API calls remain unmodified.

With this new Lithops joblib backend we can execute Scikit-learn jobs in serverless functions. In particular, we can use the Scikit-learn's Gridsearch over Lithops. The `model_selection` module that implements the `GridSearchCV` functionality has around 6650 LOC, which makes it a complex application. In this experiment, we use Gridsearch to do a hyperparameter tuning on a SGD Classifier. We used 30 MB of an Amazon Reviews dataset to perform a cross-validation of 5 folds. Each task requires a chunk of the train and test datasets. We also wanted to compare the behavior of Lithops using two different storage backends (S3 and Redis) so we used them with the default experiment settings explained previously.

The results in Fig. 11 show that with the same number of vCPUs the execution time of the VMs is between 3% and 5% lower than the execution time of Lithops with S3 and between 3% and 7% higher than the execution time of Lithops using Redis. However, for this type of problem, the scalability of Lithops follows the same progression as the VM and Lithops quickly scale to much higher levels than the VM maximum, obtaining a speedup of up to 3.6x times greater. Finally, we can see that Redis is about a 10% better than S3 when few processes are used and, from 256 processes, Redis begins to saturate but S3 continues scaling correctly. This is caused by the increasing number of concurrent reads of the training and validation data. S3 is able to serve a higher number of concurrent reads compared to Redis which is single-threaded. This implies that S3 has better bandwidth and throughput, which decreases the data read overhead, and consequently decreases the overall execution time.

#### 6.4. Proximal policy optimization

OpenAI Baselines [29] is a set of high-quality implementations of reinforcement learning algorithms. It has been open-sourced to be used as a base, around which, new ideas can be added and as a tool for comparing new approaches in the reinforcement learning field. At the time of writing the paper, the Baselines code repository contains more than 16 700 lines of Python code, so it can be

considered as a complex Python module. In this experiment, we want to verify that thanks to the access transparency provided by Lithops *multiprocessing* we can simulate the vertical scaling of a virtual machine using FaaS as processes.

We have used the multiprocessing implementation of the Proximal Policy Optimization (PPO) algorithm from OpenAI baselines. The *multiprocessing* PPO version is the second implementation released by OpenAI, and it inherits some of its structure from the first version, which was based on MPI. For that reason, the *multiprocessing* PPO uses a main-worker paradigm relying on Pipes for the main to worker communications and vice versa.

The main process is in charge of training the model (a neural network) which, for a given scenario, decides the optimum action to do in order to maximize an objective function. The worker processes are used to simulate the environment in which actions are performed and a reaction is obtained. It is important to notice that each worker process simulates an environment. The training of the model is an iterative procedure where the workers send to the main the actual state of the environment, and it responds with an action to perform in each environment. From the *multiprocessing* API, PPO uses 1 Context with the *spawn* mode by default. Associated to that Context, it creates 1 Process and 1 Pipe for each environment emulated. The communication between the main and workers (where states and actions are transmitted) is performed using the Pipe associated to each worker Process.

In this experiment, we are training a neural network to play the Atari game Breakout, which is available in the OpenAI GYM [30]. Notice that due to the TensorFlow 1 dependency, we have used Python 3.7 in this experiment.

As this algorithm requires the use of a GPU in the main process for the neural network training, the settings for this experiment have been modified. We have used an AWS p3.2xlarge VM as monolithic system, and we tried to scale it vertically using AWS Lambdas. Since the GPU is just used in the main process that runs in the Lithops orchestrator and not in the workers that just do environment simulation, the configuration of the AWS lambdas has not been modified.

The results, available in Fig. 12, show that despite the constant communication between processes and the great overhead that this entails, the combination of VM and Lithops achieves a better performance than just the VM. In more detail, the best result for the VM is achieved using 16 processes with a total execution time of 68.92 s and the best result of the VM + Lithops is achieved using 64 processes with a total execution time of 61.10 s, therefore it reduces an 11% the execution time. This validates

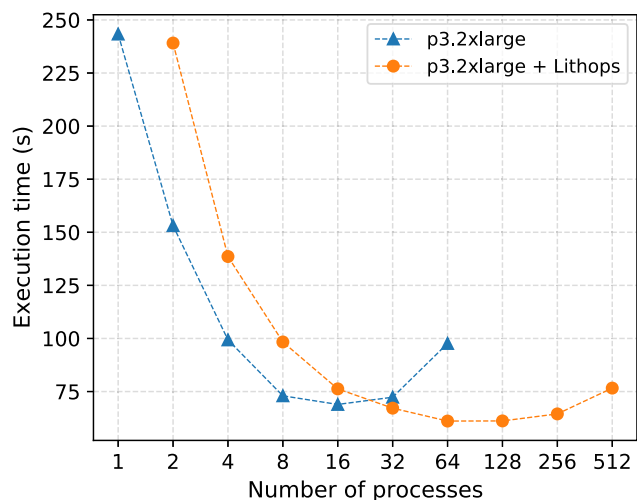


Fig. 12. PPO execution results.

that we can emulate a vertical scaling of the VM, and that it is possible to add vCPUs to a VM instantly and without prior provisioning thanks to the use of FaaS.

## 7. Insights

After studying the results of the evaluation, we have learned that we are able to transparently move local-parallel applications to distributed settings using serverless. Nevertheless, we have chosen only four representative Python applications that use the multiprocessing library. By this, we do not want to make the claim that all applications can be transparently scaled using serverless without significant degradation. In this section, we would like to discuss the insights obtained from the evaluation that have implications on the feasibility of transparency using serverless services.

### 7.1. Message passing over shared memory

In distributed systems, a message-passing model is usually used instead of a Distributed Shared Memory (DSM) model. However, there are local-parallel applications that are simply more natural to be developed with shared memory than with message passing. This presents a problem for applications that make heavy usage of local shared memory when transparently moved to a distributed environment, as their performance will be adversely affected, as seen in the experiment of Section 5.5. If the shared memory access is read-only or with infrequent light writes, it may be possible to implement optimizations that decrease the overhead penalty of accessing DSM.

### 7.2. Shared state interfaces

Clean shared memory abstractions to communicate processes are very important. Structured and consistent access to shared state requires suitable programming abstractions. In this case, Python multiprocessing design is a clear facilitator for achieving transparency. The ability to perform parallel execution in Python using threads is limited by the GIL, which prevents multiple threads from running simultaneously on multiprocessor architectures. For this reason, in Python, it is necessary to use processes to have true parallelism. Many of the principles of multiprocessing abstractions, such as Manager, are based on message passing and accessing shared objects (queues, dictionaries, lists...) instead of

traditional memory sharing. For example, in a multi-thread application written in Java, two parallel threads can access a shared object by a reference pointer. In contrast, in Python multiprocessing, two processes access to shared state by using messages through a third process (the Manager) that has the shared state. The fact that two Python processes cannot share the same address space<sup>6</sup> has facilitated the port of this library to its distributed implementation using disaggregated resources. If the code is not using adequate programming abstractions, full transparency may be impossible.

### 7.3. Latencies and overheads

Overheads are still relevant for many applications. Current Cloud settings still show relevant latency in communications, like hundreds of milliseconds to launch a serverless function, or hundred of microseconds to access in-memory storage services. We have seen that, with equal resources, the overheads generated by creating processes and by the latency of access to shared state are very noticeable. In this line, the granularity of computing tasks is clearly limited by overheads. Very fine-grained computing tasks do not make sense in the current Serverless model, since the overheads can be greater than the task run time.

### 7.4. Performance

Some parallel applications have certain advantages in Cloud Serverless settings that may help to mitigate some of the overheads.

First, Hyper-Threading may cause performance degradation in virtual machines for compute-intensive tasks using all vCPUs. Hyper-Threading makes two threads share some CPU resources like the Arithmetical Logic Unit (ALU). For computationally intensive tasks, two threads are constantly fighting for the shared resources, so the CPU cannot keep up and the execution time is degraded. In High-Performance Computing, disabling Hyper-Threading is a common practice to avoid these problems, although the capacity of effective parallelism is reduced by half. For Serverless Functions, AWS Lambda assigns a vCPU (an Hyper-Threaded CPU thread) per function with a memory configuration of 1769 MB. However, our observations show that the inefficiencies caused by Hyper-Threading in VMs do not occur in Lambda function executions. This provides an opportunity to further improve the parallelism of high-performance applications that require a full physical CPU for better performance.

Second, accessing large volumes of data in Cloud Object Storage from Serverless functions helps to aggregate bandwidth and accelerate data transfers. A single VM cannot compete with parallel data flows from multiple functions.

In addition, as we have seen in the validation of Section 6.4, disaggregated resources can serve as “accelerators” for a VM. That is, when a VM reaches the maximum occupancy of local resources, it could allocate and move computation to disaggregated resources, e.g., to serverless functions. In this way, we could benefit from both fast-access local memory for shared state-dependent processes running on the VM and high flexibility and scalability for stateless processes.

### 7.5. Fault tolerance and serverless services

The fault tolerance of our solution is based on the assumption that the underlying disaggregated resources are fault tolerant. When programming a monolithic local system, fault tolerance

<sup>6</sup> except for the multiprocessing. Array abstraction, which can only store basic C types.

**Table 5**

Application's execution time and cost for maximum speedup for AWS ECS and AWS Lambda. "Comparison" column indicates how much faster and expensive it is to execute on serverless compared to a virtual machine.

Application	VM			Lithops			Comparison	
	Exec. time	Processes	Cost	Exec. time	Processes	Cost	Cost	Time
Evolution strategies	39.11 s	96	0.044 \$	29.63 s	512	0.44 \$	9.93×	1.31×
Pandara-1el	6.36 s	96	0.007 \$	6.87 s	96	0.02 \$	2.65×	0.92×
Gridsearch	96.97 s	96	0.109 \$	28.73 s	1024	0.85 \$	7.76×	3.37×
PPO	68.92 s	16	0.058 \$	61.10 s	64	0.11 \$	2.82×	1.12×

is not taken into account because local resources do not fail. When we move to a distributed environment, if the disaggregated resources (compute, memory and storage) mask the possible failures that may occur, then the application programmer can also assume that they will not fail, and we can continue with the same programming model that does not contemplate error handling and rely on the same local programming model.

Precisely, both AWS Lambda and AWS S3 are fault-tolerant. AWS Lambda can detect and retry failed invocations, while AWS S3 objects are replicated. However, in-memory storage is still not offered as a managed service with scalability and fault tolerance. We are relying on a dedicated Redis service, which must be properly managed to ensure scalability and fault tolerance of the in-memory layer. If the data flows exceed the capacity of this intermediate node, the experiment would fail.

Regarding storage, we are now intercepting file access that is routed to Object Storage. But Object Storage has certain limitations regarding small files or read/write operations. Intensive use of such operations by applications would also preclude transparency. Serverless disaggregated memory and fine-grained storage services are needed in the Cloud.

### 7.6. Cost

Finally, we will discuss the cost of our solution. Nowadays, FaaS is more expensive than IaaS for the same resources (approximately twice as much). The reason is that they are offered as fully managed services that provide fault tolerance and automatic scaling by delegating engineering and management concerns from the user to the Cloud provider, which entails a plus in cost.

In Table 5 we see a summary of the monetary cost for the execution with the best speedup for both configurations, VM and Lithops, for all applications. The costs have been calculated using the total execution time of the application and the price per hour of the VM in AWS EC2 on demand, and the price per GB/s of AWS Lambda. In the "Comparison" column we can see a comparison of both price and time. For example, for the Evolution Strategies application, approximately ten times the cost is required for Lithops to obtain a 30% improvement in performance compared to a VM. As we can see, our solution clearly does not pay off from a purely monetary perspective.

However, we must state that the cost values only contemplate the cost for the resources used and not the time taken by the application user to manage those resources. The serverless paradigm billing model only charges for the exact resources used by the application execution, contrary to the IaaS billing model where the user is billed by allocated resources, which may be underutilized. For instance, the user may take some time between executions to analyze the results – time that is wasted by having allocated resources that are not being used at full capacity. We do not want to overemphasize on the cost aspect because it is difficult to calculate the real cost. Data analysts have to cope with an ever-growing availability of tools and services, which prove to be counterproductive because of the management burden. Approximately only a third of developers are actual Cloud experts, so many turn to serverless services because of their ease of adoption, despite the extra cost involved [31]. In addition, billing models for serverless services may unexpectedly change in the future, so the cost would have to be re-evaluated.

## 8. Conclusion

In this paper, we have demonstrated that Python's multiprocessing message-passing shared state design enables to seamlessly port local-parallel applications over disaggregated serverless resources in the Cloud. Thanks to access transparency, with just changing a single line of code, we are able to deploy a complex local-parallel application in a distributed way in the Cloud. Serverless Cloud services, such as FaaS or Object Storage, allow to massively exploit the parallelism of applications and to further reduce the execution time by increasing parallelism speedup.

We have demonstrated that applications which use stateful abstractions based on message passing, such as queues or pipes, are easy to disaggregate and that the overheads introduced are negligible, obtaining good performance in comparison to the same application running in a big standalone VM. Access transparency is a key to simplify the whole process of moving applications to the Cloud: legacy applications benefit from access transparency since architecture re-engineering would not be required anymore, and data scientists that are familiar with local-parallel programming can instantly and effortlessly scale their code in the Cloud to process bigger workloads.

Nevertheless, performance is severely affected if shared memory abstractions are heavily used, since distributed shared memory will never be as fast as local shared memory. In addition, we require programming interfaces where compute resources (such as processes) or state resources (such as queues) are clearly defined.

In conclusion, access transparency is currently possible with some caveats. However, we are optimistic that network latencies will be reduced, and therefore overheads too, so that access transparency will provide ability to program the Cloud as a parallel Super-Computer, thus hiding the complexities of distributed systems.

### CRedit authorship contribution statement

**Aitor Arjona:** Conceptualization, Software, Validation, Investigation, Writing – original draft, Visualization. **Gerard Finol:** Conceptualization, Software, Validation, Investigation, Writing – original draft, Visualization. **Pedro García López:** Conceptualization, Investigation, Writing – review & editing, Supervision, Project administration, Funding acquisition.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Data availability

Code and data is open source and available, link can be found in the article

## Acknowledgments

This work has been partially funded by the EU Horizon 2020 programme under grant agreement No. 825184 and by the Spanish Ministry of Science and Innovation, Spain and State Research Agency (Agencia Estatal de Investigación), Spain under grant agreement PID2019-106774RB-C22. Aitor Arjona is a URV Martí Franquès grant fellow.

## References

- [1] G. Coulouris, J. Dollimore, T. Kindberg, G. Blair, *Distributed systems - concepts and design* (5th edition), Addison-Wesley Longman, Inc., 2012, pp. 23–25.
- [2] A. Kalia, M. Kaminsky, D. Andersen, Datacenter RPCs can be General and Fast, in: 16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19), USENIX Association, Boston, MA, 2019, pp. 1–16, URL <https://www.usenix.org/conference/nsdi19/presentation/kalia>.
- [3] S.M. Rumble, D. Ongaro, R. Stutsman, M. Rosenblum, J.K. Ousterhout, It's time for low latency, in: *Proceedings of the 13th USENIX Conference on Hot Topics in Operating Systems, HotOS '13*, USENIX Association, USA, 2011, p. 11.
- [4] P.G. López, A. Slominski, S. Shillaker, M. Behrendt, B. Metzler, Serverless end game: Disaggregation enabling transparency, 2020, CoRR, [abs/2006.01251](https://arxiv.org/abs/2006.01251). <https://arxiv.org/abs/2006.01251>.
- [5] P.X. Gao, A. Narayan, S. Karandikar, J. Carreira, S. Han, R. Agarwal, S. Ratnasamy, S. Shenker, Network requirements for resource disaggregation, in: *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, 2016, pp. 249–264.
- [6] S. Angel, M. Nanavati, S. Sen, Disaggregation and the application, 2019, CoRR [abs/1910.13056](https://arxiv.org/abs/1910.13056). <http://arxiv.org/abs/1910.13056>.
- [7] E. Jonas, Q. Pu, S. Venkataraman, I. Stoica, B. Recht, Occupy the cloud: Distributed computing for the 99%, in: *Proceedings of the 2017 Symposium on Cloud Computing, ACM*, 2017, pp. 445–451.
- [8] V. Shankar, K. Krauth, Q. Pu, E. Jonas, S. Venkataraman, I. Stoica, B. Recht, J. Ragan-Kelley, Numpywren: serverless linear algebra, 2018, [arXiv:1810.09679](https://arxiv.org/abs/1810.09679).
- [9] S. Fouladi, R.S. Wahby, B. Shacklett, K.V. Balasubramaniam, W. Zeng, R. Bhalerao, A. Sivaraman, G. Porter, K. Winstein, Encoding, Fast and Slow: Low-Latency Video Processing Using Thousands of Tiny Threads, in: *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, USENIX Association, Boston, MA, 2017, pp. 363–376, URL <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/fouladi>.
- [10] J. Protic, M. Tomasevic, V. Milutinovic, A survey of distributed shared memory systems, in: *Proceedings of the Twenty-Eighth Annual Hawaii International Conference on System Sciences, Vol. 1*, 1995, pp. 74–84, <https://dx.doi.org/10.1109/HICSS.1995.375407>.
- [11] S.C. Kendall, J. Waldo, A. Wollrath, G. Wyant, A Note on Distributed Computing, *Tech. rep.*, Sun Microsystems, Inc., USA, 1994.
- [12] J. Sampe, M. Sanchez-Artigas, G. Vernik, I. Yehekel, P. Garcia-Lopez, Outsourcing data processing jobs with lithops, *IEEE Trans. Cloud Comput.* (2021) 1, <https://dx.doi.org/10.1109/TCC.2021.3129000>.
- [13] Y. Shan, Y. Huang, Y. Chen, Y. Zhang, LegoOS: A disseminated, distributed OS for hardware resource disaggregation, in: *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, USENIX Association, Carlsbad, CA, 2018, pp. 69–87, URL <https://www.usenix.org/conference/osdi18/presentation/shan>.
- [14] J. Zhang, Z. Ding, Y. Chen, X. Jia, B. Yu, Z. Qi, H. Guan, GiantVM: A type-II hypervisor implementing many-to-one virtualization, in: *Proceedings of the 16th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments, VEE '20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 30–44, <https://dx.doi.org/10.1145/3381052.3381324>.
- [15] J. Spillner, Transformation of python applications into function-as-a-service deployments, 2017, [arXiv:1705.08169](https://arxiv.org/abs/1705.08169).
- [16] J. Zhi, R. Wang, J. Clune, K.O. Stanley, Fiber: A platform for efficient development and distributed training for reinforcement learning and population-based methods, 2020, [arXiv:2003.11164](https://arxiv.org/abs/2003.11164).
- [17] D. Barcelona-Pons, M. Sánchez-Artigas, G. París, P. Sutra, P. García-López, On the FaaS Track: Building Stateful Distributed Applications with Serverless Architectures, in: *Proceedings of the 20th International Middleware Conference, Middleware '19*, Association for Computing Machinery, New York, NY, USA, 2019, pp. 41–54, <https://dx.doi.org/10.1145/3361525.3361535>.
- [18] W. Zhang, V. Fang, A. Panda, S. Shenker, Kappa: A programming framework for serverless computing, *SoCC '20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 328–343, <https://dx.doi.org/10.1145/3419111.3421277>.
- [19] Amazon Web Services, Developer Guide for AWS Lambda, 2021, URL <https://docs.aws.amazon.com/lambda/latest/dg/configuration-function-common.html>.
- [20] A. Wang, J. Zhang, X. Ma, A. Anwar, L. Rupprecht, D. Skourtis, V. Tarasov, F. Yan, Y. Cheng, InfiniCache: Exploiting ephemeral serverless functions to build a Cost-Effective memory cache, in: *18th USENIX Conference on File and Storage Technologies (FAST 20)*, USENIX Association, Santa Clara, CA, 2020, pp. 267–281, URL <https://www.usenix.org/conference/fast20/presentation/wang-ao>.
- [21] M. Wawrzoniak, I. Müller, R. Fraga Barcelos Paulus Bruno, G. Alonso, Boxer: Data analytics on network-enabled serverless platforms, 2021-01, <https://dx.doi.org/10.3929/ethz-b-000456492>, 11th Annual Conference on Innovative Data Systems Research (CIDR 2021); Conference Location: online; Conference Date: January 11–15, 2021; The conference lecture was held on January 12, 2021. Due to the Coronavirus (COVID-19) the conference was conducted virtually..
- [22] Amazon Web Services, Optimize CPU options, Amazon EC2 User Guide for Linux Instances, 2021, URL <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instance-optimize-cpu.html>.
- [23] A. Agache, M. Brooker, A. Iordache, A. Liguori, R. Neugebauer, P. Piwonka, D.-M. Popa, Firecracker: Lightweight virtualization for serverless applications, in: *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, USENIX Association, Santa Clara, CA, 2020, pp. 419–434, URL <https://www.usenix.org/conference/nsdi20/presentation/agache>.
- [24] Amazon Web Services, Amazon EBS volume types, 2021, URL <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-volume-types.html>.
- [25] R. Wang, J. Lehman, J. Clune, K. Stanley, Paired Open-Ended Trailblazer (POET): Endlessly Generating Increasingly Complex and Diverse Learning Environments and Their Solutions, 2019, ArXiv, [abs/1901.01753](https://arxiv.org/abs/1901.01753).
- [26] Manu NALEPA, Pandarallel, 2021, URL <https://github.com/nalepae/pandarallel>.
- [27] A. Go, R. Bhayani, L. Huang, Twitter Sentiment Classification using Distant Supervision, *Processing (2009)* 1–6, URL <http://www.stanford.edu/~alecmgo/papers/TwitterDistantSupervision09.pdf>.
- [28] Joblib developers, Joblib documentation, 2021, URL <https://joblib.readthedocs.io/en/latest/>.
- [29] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, P. Zhokhov, OpenAI baselines, 2017, <https://github.com/openai/baselines>.
- [30] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, W. Zaremba, OpenAI gym, 2016, [arXiv:arXiv:1606.01540](https://arxiv.org/abs/1606.01540).
- [31] D. Berrmbach, A. Chandra, C. Krintz, A. Gokhale, A. Slominski, L. Thamsen, E. Cavalcante, T. Guo, I. Brandic, R. Wolski, On the future of cloud engineering, in: *2021 IEEE International Conference on Cloud Engineering (IC2E)*, 2021, pp. 264–275, <https://dx.doi.org/10.1109/IC2E52221.2021.00044>.



**Aitor Arjona** is a Ph.D. student of Computer Engineering and Mathematics at the Universitat Rovira i Virgili (Spain). His research interests include cloud computing, distributed systems, and middleware. His thesis is currently focused on object storage systems and serverless computing.



**Gerard Finol** received his BSc degree in Computer Science and his BSc degree in Mathematics at University of Barcelona. He is a predoctoral researcher at Cloud and Distributed Systems Lab of University Rovira i Virgili. His research is focused on access transparency in distributed systems with functional programming paradigms.



**Pedro Garcia** is full professor of the Computer Engineering and Mathematics Department at the University Rovira i Virgili (Spain). He leads the “Cloud and Distributed Systems Lab” research group and has coordinated three large European research projects in the last years: FP7 CloudSpaces (2013–2015), H2020 IOStack (2015–2017) and H2020 CloudButton (2019–2021).