

# Towards Serverless Execution of Scientific Workflows – HyperFlow Case Study

Maciej Malawski  
AGH University of Science and Technology  
Department of Computer Science  
Krakow, Poland  
malawski@agh.edu.pl

## ABSTRACT

Scientific workflows consisting of a high number of dependent tasks represent an important class of complex scientific applications. Recently, a new type of *serverless* infrastructures has emerged, represented by such services as Google Cloud Functions or AWS Lambda. In this paper we take a look at such serverless infrastructures, which are designed mainly for processing background tasks of Web applications. We evaluate their applicability to more compute- and data-intensive scientific workflows and discuss possible ways to repurpose serverless architectures for execution of scientific workflows. A prototype workflow executor function has been developed using Google Cloud Functions and coupled with the HyperFlow workflow engine. The function can run workflow tasks on the Google infrastructure, and features such capabilities as data staging to/from Google Cloud Storage and execution of custom application binaries. We have successfully deployed and executed the Montage astronomic workflow, often used as a benchmark, and we report on initial results of performance evaluation. Our findings indicate that the simple mode of operation makes this approach easy to use, although there are costs involved in preparing portable application binaries for execution in a remote environment.

While our evaluation uses a preproduction (alpha) version of the Google Cloud Functions platform, we find the presented approach highly promising. We also discuss possible future steps related to execution of scientific workflows in serverless infrastructures, and the implications with regard to resource management for scientific applications in general.

## Keywords

Scientific workflows, cloud, serverless infrastructures

## 1. INTRODUCTION

Scientific workflows consisting of a high number of dependent tasks represent an important class of complex scientific applications that have been successfully deployed and executed in traditional cloud infrastructures, including Infrastructure as a Service (IaaS) clouds. Recently, a new type of *serverless* infrastructures emerge, represented by such services as Google Cloud Functions (GCF) [2] or AWS Lambda [1]. These services allow deployment of software in the form of functions that are executed in the provider’s infrastructure in response to specific events such as new files being uploaded to a cloud data store, messages arriving in queue systems or direct HTTP calls. This approach frees the user

from having to maintain a server, including configuration and management of virtual machines, while resource management is provided by the platform in an automated and scalable way.

In this paper we take a look at such serverless infrastructures. Although designed mainly for processing background tasks of Web applications, we nevertheless investigate whether they can be applied to more compute- and data-intensive scientific workflows. The main objectives of this paper are as follows:

- To present the main features of serverless infrastructures, comparing them to traditional infrastructure-as-a-service clouds,
- To discuss the options of using serverless infrastructures for execution of scientific workflows,
- To present our experience with a prototype implemented using HyperFlow [3] workflow engine and Google Cloud Functions (alpha version),
- To evaluate our approach using the Montage workflow [20], a real-world astronomic application,
- To discuss the costs and benefits of this approach, together with its implications for resource management of scientific workflows in emerging infrastructures.

The paper is organized as follows. We begin with an overview of serverless infrastructures in Section 2. In Section 3 we propose and discuss alternative options for serverless architectures of scientific workflow systems. Our prototype implementation, based on HyperFlow and GCF, is described in Section 4. This is followed by evaluation using the Montage application, presented in Section 5. We discuss implications for resource management in Section 6 and present related work in Section 7. Section 8 provides a summary and description of future work.

## 2. OVERVIEW OF SERVERLESS CLOUDS

Writing “serverless” applications is a recent trend, mainly addressing Web applications. It frees programmers from having to maintain a server – instead they can use a set of existing cloud services directly from their application. Examples of such services include cloud databases such as Firebase or DynamoDB, messaging systems such as Google Cloud Pub/Sub, notification services such as Amazon SNS

and so on. When there is a need to execute custom application code in the background, special “cloud functions” (hereafter simply referred to as functions) can be called. Examples of such functions are AWS Lambda and Google Cloud Functions (GCF).

Both Lambda and GCF are based on the functional programming paradigm: a function is a piece of software that can be deployed on the providers’ cloud infrastructure and it performs a single operation in response to an external event.

Functions can be triggered by:

- an event generated by the cloud infrastructure, e.g. a change in a cloud database, a file being uploaded to a cloud object store, a new item appearing in a messaging system, or an action scheduled at a specified time,
- a direct request from the application via HTTP or cloud API calls.

The cloud infrastructure which hosts the functions is responsible for automatic provisioning of resources (including CPU, memory, network and temporary storage), automatic scaling when the number of function executions varies over time, as well as monitoring and logging. The user is responsible for providing executable code in a format required by the framework. Typically, the execution environment is limited to a set of supported languages: Node.js, Java and Python in the case of AWS Lambda, and Node.js in the case of GCF. The user has no control over the execution environment, such as underlying operating system, version of the runtime libraries, etc., but can use custom libraries with package managers and even upload binary code to be executed.

Functions are thus different from Virtual Machines in IaaS clouds where the users have full control over the OS (including root access) and can customize the execution environment to their needs. On the other hand, functions free the developers from the need to configure, maintain, and manage server resources.

Cloud providers impose certain limits on the amount of resources a function can consume. In the case of AWS Lambda these limits are as follows: temporary disk space: 512 MB, number of processes and threads: 1024, maximum execution duration per request: 300 seconds. There is also a limit of 100 concurrent executions per region, but this limit can be increased on request. GCF, in its alpha version, does not specify limit thresholds. There is, however, a timeout parameter that can be provided when deploying a function and the default value is 60 seconds.

Functions are thus different from permanent and stateful services, since they are not long-running processes, but rather serve individual tasks. Resource limits indicate that such cloud functions are not currently suitable for large-scale HPC applications, but can be useful for high-throughput computing workflows consisting of many fine-grained tasks.

Functions have a fine-grained pricing model associated with them. In the case of AWS Lambda, the price is \$0.20 per 1 million requests and \$0.00001667 for every GB-second used, defined as CPU time multiplied by the amount of memory used. There are also additional charges for data transfer and storage (when DynamoDB or S3 is used). The alpha version of Google Cloud Functions does not have a public pricing policy.

Serverless infrastructures can be cost-effective compared to standard VMs. For example, the aggregate cost of running AWS Lambda functions with 1 GB memory for 1 hour is \$0.060012. This is more expensive than the t2.micro instance, which also has 1 GB of RAM but costs \$0.013 per hour. A T2.micro instance, however, offers only burstable performance, which means only a fraction of CPU time per hour is available. The smallest standard instance at AWS is m3.medium, which costs \$0.067 per hour, but gives 3.75 GB of RAM. Cloud functions are thus more suitable for variable load conditions while standard instances can be more economical for applications with stable workloads.

### 3. OPTIONS FOR EXECUTION OF SCIENTIFIC WORKFLOWS IN SERVERLESS INFRASTRUCTURES

In light of the identified features and limitations of serverless infrastructures and cloud functions, we can discuss the option of using them for execution of scientific workflows. We will start with a traditional execution model in IaaS cloud with no cloud functions (1), then present the queue model (2), direct executor model (3), bridge model (4), and decentralized model (5). These options are schematically depicted in Fig. 1, and discussed in detail further on.

#### 3.1 Traditional model

The traditional model assumes the workflow is running in a standard IaaS cloud. In this model the workflow execution follows the well-known master-worker architecture, where the master node runs a workflow engine, tasks that are ready for execution are submitted to a queue, and worker nodes process these tasks in parallel when possible. The master node can be deployed in the cloud or outside of the cloud, while worker nodes are usually deployed as VMs in a cloud infrastructure. The worker pool is typically created on demand and can be dynamically scaled up or down depending on resource requirements.

Such a model is represented e.g. by Pegasus and HyperFlow. The Pegasus Workflow Management System [6] uses HTCondor [18] to maintain its queue and manage workers. HyperFlow [3] is a lightweight workflow engine based on Node.js – it uses RabbitMQ as its queue and AMQP Executors on worker nodes. The deployment options of HyperFlow on grids and clouds are discussed in detail in [4].

In this model the user is responsible for management of resources comprising the worker pool. The pool can be provisioned statically, which is commonly done in practice, but there is also ongoing research on automatic or dynamic resource provisioning for workflow applications [13, 15], which is a non-trivial task.

In the traditional cloud workflow processing model there is a need for some storage service to store input, output and temporary data. There are multiple options for data sharing [9], but one of the most widely used approaches is to rely on existing cloud storage, such as Amazon S3 or Google Cloud Storage. This option has the advantage of providing a permanent store so that data is not lost after the workflow execution is complete and the VMs are terminated.

#### 3.2 Queue model

This model is similar to the traditional model: the master node and the queue remain unchanged, but the worker is

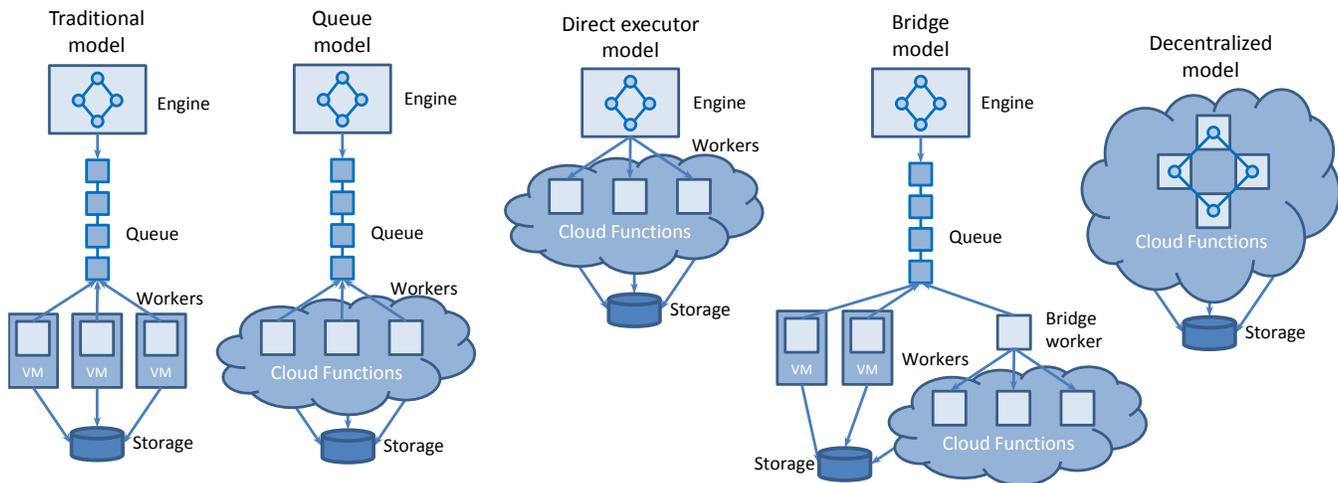


Figure 1: Options of serverless architectures for execution of scientific workflows.

replaced by a cloud function. Instead of running a pool of VMs with workers a set of cloud functions is spawned. Each function fetches a task from a queue and processes it, returning results via the queue.

The main advantage of this model is its simplicity, since it only requires changes in the worker module. This may be simple if the queue uses a standard protocol, such as AMQP in the case of HyperFlow Executor, but in the case of Pegasus and HTCondor a Condor daemon (`condor_startd`) must run on the worker node and communicate using a proprietary Condor protocol. In this scenario implementing a worker as a cloud function would require more effort.

Another advantage of the presented model is the ability to combine the workers implemented as functions with other workers running e.g. in a local cluster or in a traditional cloud. This would also enable concurrent usage of cloud functions from multiple providers (e.g. AWS and Google) when such a multi-cloud scenario is required.

An important issue associated with the queue model is how to trigger the execution of the functions. If a native implementation of the queue is used (e.g. RabbitMQ as in HyperFlow), it is necessary to trigger a function for each task added to the queue. This can be done by the workflow engine or by a dedicated queue monitoring service. Other options include periodic function execution or recursive execution: a function can itself trigger other functions once it finishes processing data.

To ensure a clean serverless architecture another option is to implement the queue using a native cloud service which is already integrated with cloud functions. In the case of AWS Lambda one could implement the queue using DynamoDB: here, a function could be triggered by adding a new item to a task table. In the case of GFC, a Google Cloud Pub/Sub service can be used for the same purpose. Such a solution, however, would require more changes in the workflow engine and would not be easy to deploy in multi-cloud scenarios.

### 3.3 Direct executor model

This is the simplest model and requires only a workflow engine and a cloud function that serves as a task executor. It eliminates the need for a queue since the workflow engine can trigger the cloud function directly via API/HTTP

calls. Regarding development effort, it requires changes in the master and a new implementation of the worker.

An advantage of this model is its cleanliness and simplicity, but these come at the cost of tight master-worker coupling. Accordingly, it becomes more difficult to implement the multi-cloud scenario, since the workflow engine would need to be able to dispatch tasks to multiple cloud function providers.

### 3.4 Bridge model

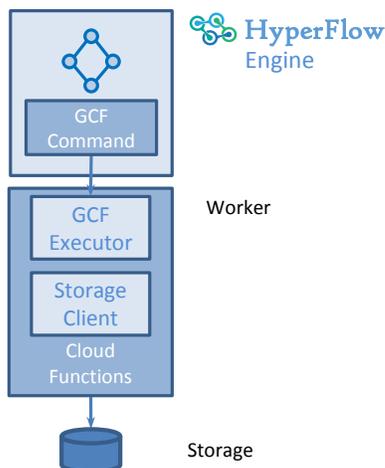
This solution is more complex but it preserves the decoupling of the master from the worker, using a queue. In this case the master and the queue remain unchanged, but a new type of bridge worker is added. It fetches tasks from the queue and dispatches them to the cloud functions. Such a worker needs to run as a separate service (daemon) and can trigger cloud functions using the provider-specific API.

The decoupling of the master from the worker allows for more complex and flexible scenarios, including multi-cloud deployments. A set of bridge workers can be spawned, each dispatching tasks to a different cloud function provider. Moreover, a pool of workers running in external distributed platforms, such as third-party clouds or clusters, can be used together with cloud functions.

### 3.5 Decentralized model

This model re-implements the whole workflow engine in a distributed way using cloud functions. Each task of a workflow is processed by a separate function. These functions can be triggered by (a) new data items uploaded to cloud storage, or (b) other cloud functions, i.e. predecessor tasks triggering their successor tasks following completion. Option (a) can be used to represent data dependencies in a workflow while option (b) can be used to represent control dependencies.

In the decentralized model the structure and state of workflow execution have to be preserved in the system. The system can be implemented in a fully distributed way, by deploying a unique function for each task in the workflow. In this way, the workflow structure is mapped to a set of functions and the execution state propagates by functions being triggered by their predecessors. Another option is to deploy



**Figure 2: Architecture of the prototype based on HyperFlow and Google Cloud Functions**

a generic task executor function and maintain the workflow state in a database, possibly one provided as a cloud service.

The advantages of the decentralized approach include fully distributed and serverless execution, without the need to maintain a workflow engine. The required development effort is extensive, since it requires re-implementation of the whole workflow engine. A detailed design of such an engine is out of scope of this paper, but remains an interesting subject of future research.

### 3.6 Summary of options

As we can see, cloud functions provide multiple integration options with scientific workflow engines. The users need to decide which option is best for them based on their requirements, most notably the allowed level of coupling between the workflow engine and the infrastructure and the need to run hybrid or cross-cloud deployments where resources from more than one provider are used in parallel. We consider the fully decentralized option as an interesting future research direction, while in the following sections we will focus on our experience with a prototype implemented using the direct executor model.

## 4. PROTOTYPE BASED ON HYPERFLOW

To evaluate the feasibility of our approach we decided to develop a prototype using the HyperFlow engine and Google Cloud Functions, applying the direct executor model. This decision has several reasons. First, HyperFlow is implemented in Node.js, while GCF supports Node.js as a native function execution environment. This good match simplifies development and debugging, which is always non-trivial in a distributed environment. Our selection of the direct execution model was motivated by the extensible design of HyperFlow, which can associate with each task in a workflow a specific executor function responsible for handling command-line tasks. Since GCF provides a direct triggering mechanism of cloud functions using HTTP calls, we can apply existing HTTP client libraries for Node.js, plugging support for GCF into HyperFlow as a natural extension.

### 4.1 Architecture and components

The schematic diagram of the prototype is shown in Fig. 2. The HyperFlow engine is extended with the `GCFCommand` function which is responsible for communication with GCF. It is a replacement for `AMQPCommand` function, which is used in the standard HyperFlow distributed deployment with AMQP protocol and RabbitMQ. `GCFCommand` sends the task description in a JSON-encoded message to the cloud function. The `GCF Executor` is the main cloud function which needs to be deployed on the GCF platform. It processes the message, and uses the `Storage Client` for staging in and out the input and output data. It uses Google Cloud Storage and requests parallel transfers to speed up download and upload of data. `GCF Executor` calls the executable which needs to be deployed together with the function. GCF supports running own Linux-based custom binaries, but the user has to make sure that the binary is portable, e.g. by statically linking all of its dependencies. Our architecture is thus purely serverless, with the HyperFlow engine running on a client machine and directly relying only on cloud services such as GCF and Cloud Storage.

### 4.2 Fault tolerance

Transient failures are a common risk in cloud environments. Since execution of a possibly large volume of concurrent HTTP requests in a distributed environment is always prone to errors caused by various layers of network and middleware stacks, the execution engine needs to be able to handle such failures gracefully and attempt to retry failed requests.

In the case of HyperFlow, the Node.js ecosystem appears very helpful in this context. We used the `requestretry` library for implementing the HTTP client, which allows for automatic retry of failed requests with a configurable number of retries (default: 5) and delay between retries (default: 5 seconds). Our prototype uses these default settings, but in the future it will be possible to explore more advanced error handling policies taking into account error types and patterns.

## 5. EVALUATION USING MONTAGE WORKFLOW

Based on our prototype which combines HyperFlow and Google Cloud Functions, we performed several experiments to evaluate our approach. The goals of the evaluation are as follows:

- To validate the feasibility of our approach, i.e. to determine whether it is practical to execute scientific workflows in serverless infrastructures.
- To measure performance characteristics of the execution environment in order to provide hints for resource management.

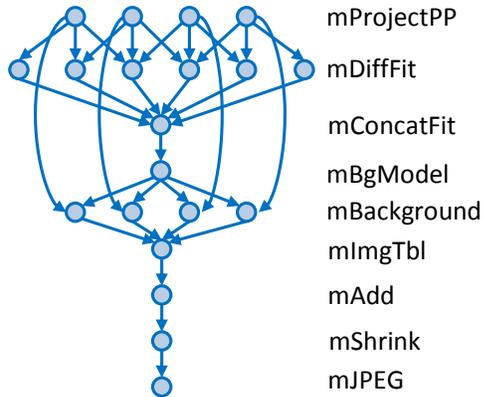
Details regarding our sample application, experiment setup and results are provided below.

### 5.1 Montage workflow and experiment setup

*Montage application.*

For our study we selected the Montage [8] application, which is an astronomic workflow. It is often used for various

benchmarks and performance evaluation, since it is open-source and has been widely studied by the research community. The application processes a set of input images from astronomic sky surveys and constructs a single large-scale mosaic image. The structure of the workflow is shown in Fig. 3: it consists of several stages which include parallel processing sections, reduction operations and sequential processing.



**Figure 3: Structure of the Montage workflow used for experiments**

The size of the workflow, i.e. the number of tasks, depends on the size of the area of the target image, which is measured in angular degrees. For example, a small-scale Montage workflow consists of 43 tasks, with 10 parallel mProjectPP tasks and 17 mDiffFit tasks, while more complex workflows can involve thousands of tasks. In our experiments we used the Montage 0.25 workflow with 43 tasks, and the Montage 0.4 workflow with 107 tasks.

### Experiment setup.

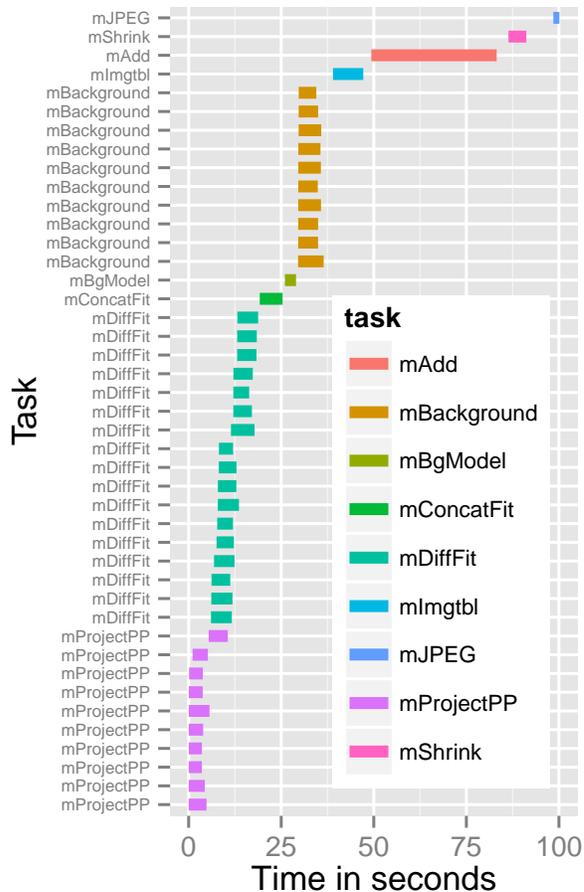
We used a recent version of HyperFlow and an Alpha version of Google Cloud Functions. The HyperFlow engine was installed on a client machine with Ubuntu 14.04 LTS Linux and Node.js 4.5.0. For staging the input and output data, as well as for temporary storage, we used a Google Cloud Storage bucket with standard options. Both Cloud Functions and Cloud Storage were located in the `us-central-1` region, while the client machine was located in Europe.

### Data preparation and handling.

To run the Montage workflow in our experiments all input data needs to be uploaded to the cloud storage first. For each workflow run, a separate subfolder in the Cloud Storage bucket is created. The subfolder is then used for exchange of intermediate data and for storing the final results. Data can be conveniently uploaded using a command-line tool which supports parallel transfers. The web-based Google Cloud Console is useful for browsing results and displaying the resulting JPEG images.

## 5.2 Feasibility

To assess the feasibility of our approach we tested our prototype using the Montage 0.25 workflow. We collected task execution start and finish timestamps, which give the total duration of cloud function execution. This execution time also includes data transfers. Based on the collected



**Figure 4: Sample run of Montage 0.25 workflow.**

execution traces we plotted Gantt charts. Altogether, several runs were performed and an example execution trace (representative of all runs) is shown in Fig. 4.

Montage 0.25 is a relatively small-scale workflow, but the resulting plot clearly reveals that the cloud function-based approach works well in this case. We can observe that the parallel tasks of the workflow (mProjectPP, mDiffFit and mBackground) are indeed short-running and can be processed in parallel. The user has no control over the level of parallelism, but the cloud platform is able to process tasks in a scalable way, as stated in the documentation. We also observe no significant delays between task execution and can attribute this to the fact that the requests between HyperFlow engine and the cloud functions are transmitted using HTTP over a wide-area network, including a trans-Atlantic connection.

Similar results were obtained for the Montage 0.4 workflow which consists of 107 tasks, however the corresponding detailed plots are not reproduced here for reasons of readability. It should be noted that while the parallel tasks of Montage are relatively fine-grained, the execution time of sequential processing tasks such as mImgTbl and mAdd grows with the size of the workflow and can exceed the default limit of 60 seconds imposed upon cloud function execution. This limit can be extended when deploying the cloud function,

but there is currently no information regarding the maximum duration of such requests. We can only expect that such limits will increase as the platforms become more mature. This was indeed the case with Google App Engine, where the initial request limit was increased from 30 seconds to 10 minutes [14].

### 5.3 Deployment size and portability

Our current approach requires us to deploy the cloud function together with all the application binaries. The Google Cloud Function execution environment enables inclusion of dependencies in Node.js libraries packaged using Node Package Manager (NPM) which are automatically installed when the function is deployed. Moreover, the user can provide a set of JavaScript source files, configuration and binary dependencies to be uploaded together with the function.

In the case of the Montage application, the users need to prepare application binaries in a portable format. Since Montage is distributed in source<sup>1</sup> format, it can be compiled and statically linked with all libraries making it portable to any Linux distribution.

The size of the Montage binaries is 50 MB in total, and 20 MB in a compressed format, which is used for deployment. We consider this deployment size as practical in most cases. We should note that deployment of the function is performed only once, prior to workflow execution. Of course, when the execution environment needs to instantiate the function or create multiple instances for scale-out scenarios, the size of each instance may affect performance, so users should try to minimize the volume of the deployment package. It is also worth noting that such binary distributions are usually more compact than the full images of virtual machines used in traditional IaaS clouds. Unfortunately, if the the source distribution or portable binary is not available, then it may not be possible to deploy it as a cloud function. One useful option would be to allow deployment of container-based images, such as Docker images, but this is currently not supported.

### 5.4 Variability

Variability is an important metric of cloud infrastructures, since distribution and resource sharing often hamper consistent performance. To measure the variability of GCF while executing scientific workflows, we collected the duration of parallel task execution in the Montage (0.25 degree) workflow – specifically, mBackground, mDiffFit and mProjectPP – running 10 workflow instances over a period of one day.

Results are shown in Fig. 5. We can see that the distribution of tasks is moderately wide, with the inter-quartile range of about 1 second width. The distribution is skewed towards longer execution times, up to 7 seconds, while the median is about 4 seconds. It is important that we do not observe any significant outliers. We have to note that the execution times of the tasks themselves vary (they are not identical) and that the task duration includes data transfer to/from cloud storage. Having taken this into account we can conclude that the execution environment behaves consistently in terms of performance, since the observed variation is rather low. Further studies and long-term monitoring would be required to determine whether such consistency is preserved over time.

<sup>1</sup><http://montage.ipac.caltech.edu/>

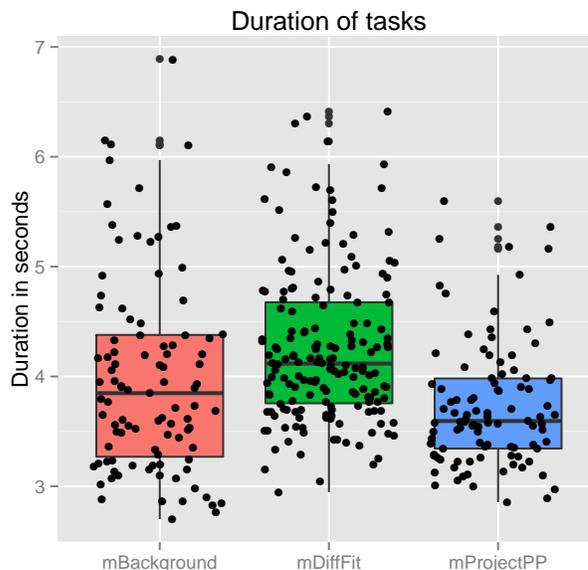


Figure 5: Distribution of execution times of parallel tasks of the Montage 0.25 workflow.

## 6. DISCUSSION

The experiments conducted with our prototype implementation confirm the feasibility of our approach to execution of scientific workflows in serverless infrastructures. There are, however, some limitations that need to be emphasized here, and some interesting implications for resource management of scientific workflows in such infrastructures.

### 6.1 Granularity of tasks

Granularity of tasks is a crucial issue which determines whether a given application is well suited for processing using serverless infrastructures. It is obvious that for long-running HPC applications a dedicated supercomputer is a better option. On the other hand, for high-throughput computing workloads distributed infrastructures such as grids and clouds have proven useful. Serverless infrastructures can be considered similar to these high-throughput infrastructures, but they usually have shorter limits of task execution size (300 seconds in the case of AWS Lambda, 60-second default timeout for GCF). While these limits may vary, may be configurable or may change over time, we must assume that each infrastructure will always impose some kind of limit, which will constrain the types of supported workflows to those consisting of relatively fine-grained tasks. Many high-throughput workflows can fit into these constraints, but for the rest, other solutions should be developed, such as hybrid approaches.

### 6.2 Hybrid solutions

In addition to purely serverless solutions, one can propose hybrid approaches, such as the one outlined in Section 3. The presented bridge model is a typical hybrid solution which combines traditional VMs with cloud functions for lightweight tasks. This architecture can overcome the limitations of cloud functions, such as the need to create

custom binaries or execution time limits.

The hybrid approach can also be used to minimize costs and optimize throughput. Such optimization should be based on cost analysis of leasing a VM and calling a cloud function, assuming that longer-term lease of resources typically corresponds to lower unit costs. This idea is generally applicable to hybrid cloud solutions [21]. For example, it may be more economical to lease VMs for long-running sequential parts of the workflow and trigger cloud functions for parallel stages, where spawning VMs that are billed on an hourly basis would be more costly. It may also prove interesting to combine cloud functions with spot instances or burstable [11] instances, which are cheaper but have varying performance and reliability characteristics.

The hybrid approach can also help resolve issues caused by the statelessness and transience of cloud functions, where no local data is preserved between function calls. By adding a traditional VM as one of the executor units, data transfers can be significantly reduced in the case of tasks that need to access to the same set of data multiple times.

### 6.3 Resource management and autoscaling

The core idea behind serverless infrastructures is that they free the users from having to manage the server – and this also extends to clusters of servers. Decisions concerning resource management and autoscaling are thus made by the platform based on the current workload, history, etc. This is useful for typical Web or mobile applications that have interactive usage patterns and whose workload depends on user behavior. With regard to scientific workflows which have a well-defined structure, there is ongoing research on scheduling algorithms for clusters, grids and clouds. The goal of these algorithms is to optimize such criteria as time or cost of workflow execution, assuming that the user has some control over the infrastructure. In the case of serverless infrastructures the user does not have any control over the execution environment. The providers would need to change this policy by adding more control or the ability to specify user preferences regarding the performance.

For example, users could specify priorities when deploying cloud functions, and a higher priority would mean faster response time, quicker autoscaling, etc., but at an additional price. Lower-priority functions could have longer execution times, possibly relying on resource scavenging, but at a lower cost. Another option would be to allow users to provide hints regarding expected execution times or anticipated parallelism level. Such information could be useful for internal resource managers to better optimize the execution environment and prepare for demand spikes, e.g. when many parallel tasks are launched by a workflow.

Adding support for cooperation between the application and the internal resource manager of the cloud platform would open an interesting area for research and optimization of applications and infrastructures which both users and providers could potentially benefit from.

## 7. RELATED WORK

Although scientific workflows in clouds have been widely studied, research focus is typically on IaaS and is little related work regarding serverless or other alternative types of infrastructures.

An example of using AWS Lambda for analyzing genomics data comes from the AWS blog [5]. The authors show how

to use R, AWS Lambda and the AWS API gateway to process a large number of tasks. Their use case is to compute some statistics for every gene in the genome, which gives about 20,000 tasks in an embarrassingly parallel problem. This work is similar to ours, but our approach is more general, since we show how to implement generic support for scientific workflows.

A detailed performance and cost comparison of traditional clouds with microservices and the AWS Lambda serverless architecture is presented in [19]. An enterprise application was benchmarked and results show that serverless infrastructures can introduce significant savings without impacting performance. Similarly, in [20] the authors discuss the advantages of using cloud services and AWS Lambda for systems that require higher resilience. They show how serverless infrastructures can reduce costs in comparison to traditional IaaS resources and the spot market. Although these use cases are different from our scientific scenario, we believe that serverless infrastructures offer an interesting option for scientific workflows.

An interesting general discussion on the economics of hybrid clouds is presented in [21]. The author shows that even if when a private cloud is strictly cheaper (per unit) than public clouds, a hybrid solution can result in a lower overall cost in the case of a variable workload. We expect that a similar effect can be observed in the case of a hybrid solution combining traditional and serverless infrastructures for scientific applications which often have a wide range of granularity of tasks.

Regarding the use of alternative cloud solutions for scientific applications, there is work on evaluation of Google App Engine for scientific applications [16, 14]. Google App Engine is a Platform-as-a-Service cloud, designed mostly for Web applications, but with additional support for processing of background tasks. App Engine can be used for running parameter-study high-throughput computing workloads, and there are similar task processing time limits as in the case of serverless infrastructures. The difference is that the execution environment is more constrained, e.g. only one application framework is allowed (such as Java or Python) and there is no support for native code and access to local disk. For these reasons, we consider cloud functions such as AWS Lambda or Google Cloud Functions as a more interesting option for scientific applications.

The concept of cloud functions can be considered as an evolution of former remote procedure call concepts, such as GridRPC [17], proposed and standardized for Grid computing. The difference between these solutions and current cloud functions is that the latter are supported by commercial cloud providers with emphasis on ease of use and development productivity. Moreover, the granularity of tasks processed by current cloud functions tends to be finer, so we need to follow the development of these technologies to further assess their applicability to scientific workflows.

A recently developed approach to decentralized workflow execution in clouds is represented by Flowbster [10], which also aims at serverless infrastructures. We can expect that more such solutions will emerge in the near future.

The architectural concepts of scientific workflows are discussed in the context of component and service architectures [7]. Cloud functions can be considered as a specific class of services or components, which are stateless and can be deployed in cloud infrastructures. They do not impose

any rules of composition, giving more freedom to developers. The most important distinction is that they are backed by the cloud infrastructure which is responsible for automatic resource provisioning and scaling.

The architectures of cloud workflow systems are also discussed in [12]. We believe that such architectures need to be re-examined as new serverless infrastructures become more widespread.

Based on the discussion of related work we conclude that our paper is likely the first attempt to use serverless clouds for scientific workflows and we expect that more research in this area will be needed as platforms become more mature.

## 8. SUMMARY AND FUTURE WORK

In this paper we have presented our approach to combining scientific workflows with the emerging serverless clouds. We believe that such infrastructures based on the concept of cloud functions, such as AWS Lambda or Google Cloud Functions, provide an interesting alternative not only for typical enterprise applications, but also for scientific workflows. We have discussed several options for designing serverless workflow execution architectures, including queue-based, direct executor, hybrid (bridged) and decentralized ones.

To evaluate the feasibility of our approach we implemented a prototype based on the HyperFlow engine and Google Cloud Functions, and evaluated it with the real-world Montage application. Experiments with small-scale workflows consisting of 43 and 107 tasks confirm that the GCF platform can be successfully used, and that it does not introduce significant delays. We have to note that the application needs to be prepared in a portable way to facilitate execution on such infrastructure and that this may be an issue for more complex scientific software packages.

Our paper also presents some implications of serverless infrastructures for resource management of scientific workflows. First, we observed that not all workloads are suitable due to execution time limits, e.g. 5 minutes in the case of AWS Lambda – accordingly, the granularity of tasks has to be taken into account. Next, we discussed how hybrid solutions combining serverless and traditional infrastructures can help optimize the performance and cost of scientific workflows. We also suggest that adding more control or the ability to provide priorities or hints to cloud platforms could benefit both providers and users in terms of optimizing performance and cost.

Since this is a fairly new topic, we see many options for future work. Further implementation work on development and evaluation of various serverless architectures for scientific workflows is needed, with the decentralized option regarded as the greatest challenge. A more detailed performance evaluation of different classes of applications on various emerging infrastructures would also prove useful to better understand the possibilities and limitations of this approach. Finally, interesting research can be conducted in the field of resource management for scientific workflows, to design strategies and algorithms for optimizing time or cost of workflow execution in the emerging serverless clouds.

## Acknowledgments

This work is partially supported by the National Centre for Research and Development (NCBiR), Poland, project PBS1/B9/18/2013. AGH grant no. 11.11.230.124 is also

acknowledged. The author would like to thank the Google Cloud Functions team for the opportunity to use the alpha version of their service.

## 9. REFERENCES

- [1] AWS Lambda - Serverless Compute, 2016. <https://aws.amazon.com/lambda/>.
- [2] Cloud Functions - Serverless Microservices | Google Cloud Platform, 2016. <https://cloud.google.com/functions/>.
- [3] B. Balis. HyperFlow: A model of computation, programming approach and enactment engine for complex distributed workflows. *Future Generation Computer Systems*, 55:147–162, sep 2016.
- [4] B. Balis, K. Figiela, M. Malawski, M. Pawlik, and M. Bubak. A Lightweight Approach for Deployment of Scientific Workflows in Cloud Infrastructures. In R. Wyrzykowski, E. Deelman, J. Dongarra, K. Karczewski, J. Kitowski, and K. Wiatr, editors, *Parallel Processing and Applied Mathematics: 11th International Conference, PPAM 2015, Krakow, Poland, September 6-9, 2015. Revised Selected Papers, Part I*, pages 281–290, Cham, 2016. Springer International Publishing.
- [5] Bryan Liston. Analyzing Genomics Data at Scale using R, AWS Lambda, and Amazon API Gateway | AWS Compute Blog, 2016. <http://tinyurl.com/h7vyboo>.
- [6] E. Deelman, K. Vahi, G. Juve, M. Rynge, S. Callaghan, P. J. Maechling, R. Mayani, W. Chen, R. Ferreira da Silva, M. Livny, and K. Wenger. Pegasus, a workflow management system for science automation. *Future Generation Computer Systems*, 46:17–35, may 2015.
- [7] D. Gannon. *Component Architectures and Services: From Application Construction to Scientific Workflows*, pages 174–189. Springer London, London, 2007.
- [8] J. C. Jacob, D. S. Katz, G. B. Berriman, J. C. Good, A. Laity, E. Deelman, C. Kesselman, G. Singh, M.-H. Su, T. Prince, and Others. Montage: a grid portal and software toolkit for science-grade astronomical image mosaicking. *International Journal of Computational Science and Engineering*, 4(2):73–87, 2009.
- [9] G. Juve, E. Deelman, K. Vahi, G. Mehta, B. Berriman, B. P. Berman, and P. Maechling. Data Sharing Options for Scientific Workflows on Amazon EC2. In *SC '10 Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '10*, pages 1–9. IEEE Computer Society, 2010.
- [10] P. Kacsuk, J. Kovacs, and Z. Farkas. Flowbster: Dynamic creation of data pipelines in clouds. In *Digital Infrastructures for Research event, Krakow, Poland, 28-30 September 2016*, 2016.
- [11] P. Leitner and J. Scheuner. Bursting with Possibilities – An Empirical Study of Credit-Based Bursting Cloud Instance Types, dec 2015.
- [12] X. Liu, D. Yuan, G. Zhang, W. Li, D. Cao, Q. He, J. Chen, and Y. Yang. *The Design of Cloud Workflow Systems*. Springer New York, New York, NY, 2012.
- [13] M. Malawski, G. Juve, E. Deelman, and J. Nabrzyski. Algorithms for cost-and deadline-constrained

- provisioning for scientific workflow ensembles in IaaS clouds. *Future Generation Computer Systems*, 48:1–18, 2015.
- [14] M. Malawski, M. Kuzniar, P. Wojcik, and M. Bubak. How to Use Google App Engine for Free Computing. *IEEE Internet Computing*, 17(1):50–59, Jan 2013.
- [15] M. Mao and M. Humphrey. Auto-scaling to minimize cost and meet application deadlines in cloud workflows. In *SC '11 Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '11, Seattle, Washington, 2011. ACM.
- [16] R. Prodan, M. Sperk, and S. Ostermann. Evaluating High-Performance Computing on Google App Engine. *IEEE Software*, 29(2):52–58, Mar 2012.
- [17] K. Seymour, H. Nakada, S. Matsuoka, J. Dongarra, C. Lee, and H. Casanova. Overview of gridrpc: A remote procedure call api for grid computing. In *International Workshop on Grid Computing*, pages 274–278. Springer, 2002.
- [18] D. Thain, T. Tannenbaum, and M. Livny. Distributed computing in practice: the Condor experience. *Concurrency and Computation: Practice and Experience*, 17(2-4):323–356, 2005.
- [19] M. Villamizar, O. Garces, L. Ochoa, H. Castro, L. Salamanca, M. Verano, R. Casallas, S. Gil, C. Valencia, A. Zambrano, and M. Lang. Infrastructure cost comparison of running web applications in the cloud using aws lambda and monolithic and microservice architectures. In *2016 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, pages 179–182, May 2016.
- [20] B. Wagner and A. Sood. Economics of Resilient Cloud Services. In *1st IEEE International Workshop on Cyber Resilience Economics*, Aug 2016.
- [21] J. Weinman. Hybrid Cloud Economics. *IEEE Cloud Computing*, 3(1):18–22, Jan 2016.