



HAL
open science

Performance Model of MapReduce Iterative Applications for Hybrid Cloud Bursting

Francisco Clemente-Castello, Bogdan Nicolae, Rafael Mayo, Juan Carlos
Fernandez

► **To cite this version:**

Francisco Clemente-Castello, Bogdan Nicolae, Rafael Mayo, Juan Carlos Fernandez. Performance Model of MapReduce Iterative Applications for Hybrid Cloud Bursting. IEEE Transactions on Parallel and Distributed Systems, 2018, 29 (8), pp.1794-1807. hal-01999033

HAL Id: hal-01999033

<https://hal.science/hal-01999033>

Submitted on 30 Jan 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Performance Model of MapReduce Iterative Applications for Hybrid Cloud Bursting

Francisco J. Clemente-Castelló*, Bogdan Nicolae†,
Rafael Mayo*, Juan Carlos Fernández*

*Universitat Jaume I, Spain

Email: {fclement, mayo, jfernand}@uji.es

†Argonne National Laboratory, USA

Email: bogdan.nicolae@acm.org



Abstract—Hybrid cloud bursting (i.e., leasing temporary off-premise cloud resources to boost the overall capacity during peak utilization) can be a cost-effective way to deal with the increasing complexity of big data analytics, especially for iterative applications. However, the low throughput, high latency network link between the on-premise and off-premise resources (“weak link”) makes maintaining scalability difficult. While several data locality techniques have been designed for big data bursting on hybrid clouds, their effectiveness is difficult to estimate in advance. Yet such estimations are critical, because they help users decide whether the extra pay-as-you-go cost incurred by using the off-premise resources justifies the runtime speed-up. To this end, the current paper presents a performance model and methodology to estimate the runtime of iterative MapReduce applications in a hybrid cloud-bursting scenario. The paper focuses on the overhead incurred by the weak link at fine granularity, for both the map and the reduce phases. This approach enables high estimation accuracy, as demonstrated by extensive experiments at scale using a mix of real-world iterative MapReduce applications from standard big data benchmarking suites that cover a broad spectrum of data patterns. Not only are the produced estimations accurate in absolute terms compared with experimental results, but they are also up to an order of magnitude more accurate than applying state-of-art estimation approaches originally designed for single-site MapReduce deployments.

Index Terms—Hybrid Cloud; Big Data Analytics; Iterative Applications; MapReduce; Performance Prediction; Runtime Estimation

1 INTRODUCTION

An important class of problems running on private clouds is *big data analytics*. However, with data sizes exploding (zettabytes predicted by 2020 [1]) and applications becoming increasingly complex, private clouds struggle to accommodate the required scale and scope. Often there is simply not enough capacity to run the desired analytics, or it is difficult to obtain the desired results within a given deadline. In addition, the rich, shared big data ecosystem facilitated by public cloud computing (large amounts of data exploitable from multiple data sources and users) opens many new opportunities for combined analytics that potentially enable new insight beyond what is possible within the scope of

a private cloud alone. In this context, *cloud bursting* [2] has seen a rapid increase in popularity among big data analytics users. It is a form of *hybrid cloud computing* that enables temporary boosting of on-premise resources managed by a private cloud with additional off-premise resources from a public cloud provider, for the purpose of overcoming the limitations of private data centers only when necessary (e.g., during peak utilization) in a flexible, cost-efficient pay-as-you-go fashion.

Enabling cloud bursting for big data analytics at large scale poses a major challenge: unlike conventional datacenters where big data analytics applications and middleware run on top of physically colocated IT resources with high-speed interconnections, the use of both on-premise and off-premise resources is bottlenecked by a “weak link” between them that is often orders of magnitude slower (e.g., 1 Gbps links between on-premise virtual machines but a shared 100 Mbps link to access off-premise VMs from a public cloud). This effect has multiple implications at the level of the runtime and storage layer, prompting the need for new “hybrid cloud big data analytics” approaches.

In our previous work [3], [4], we showed that *iterative* applications are particularly well positioned to take advantage of hybrid clouds. This class of applications typically reuses a large amount of invariant input data. Hence, the runtime overhead introduced by shipping data off-premise over the weak link can be offset by leveraging locality awareness to reuse it at each iteration. Thus, by taking advantage of these properties, we introduced two complementary techniques to accelerate iterative big data applications on hybrid clouds, which we illustrated based on Hadoop MapReduce [5]: (1) extended off-premise HDFS storage layer using asynchronous rack-aware replica rebalancing and (2) locality-enforced scheduling that avoids redundant data transfers over the weak link. Thanks to these techniques, we have shown significant speed-up compared with the default Hadoop implementation designed for a single data-center [6]. In addition, we showed that such techniques can perform close to the lower bound; that is,

the performance is close to the case when more on-premise resources are added to match the capacity of the hybrid setup.

Despite the potential to achieve a significant speed-up, however, the extra pay-as-you-go cost for the temporary provisioning of the off-premise virtual machines (VMs) can be significant. Thus, from a business perspective, the ability to *estimate the speed-up in advance* for a given hybrid configuration is critical, in order to be able to decide whether the return of the investment is justified *before committing* to any extra pay-as-you-go costs.

We focus here on this challenge. Specifically, we extend our previous work with a generic performance model that can be used to estimate the runtime of iterative MapReduce applications in a hybrid cloud. A key advantage of the model is that it offers both a pessimistic (upper) and optimistic (lower) runtime estimation, which enable more informed decision-making that is adapted to the inherent performance fluctuations exhibited by the off-premise VMs due to multitenancy on the public clouds. We summarize our contributions as follows.

- We elaborate on the fundamental issues in operating iterative MapReduce applications over hybrid cloud setups comprising both on- and off-premise VMs. In particular, we further develop the main issue VMs. In particular, we further develop the main issue discussed in our previous work (lacking data locality on the off-premise part and associated consequences) from multiple angles: I/O interactions with the underlying storage layer, task scheduling, and data shuffling (Section 3).
- We propose a methodology that combines analytical modeling with synthetic benchmarking to estimate the time to solution in a hybrid setup, including all fine-grained overhead associated with the map phase and the reduce phase (shuffle, sort, reduce). This model extends our preliminary efforts in this direction [4] that addressed the map phase only (Section 4).
- We evaluate our approach in a series of experiments that involve four representative real-life iterative MapReduce applications from standardized big data benchmarks that cover a broad range of use cases. Our experiments demonstrate small errors between the runtime estimations and the actual measured values, which are up to an order of magnitude smaller than using state-of-art MapReduce runtime estimation approaches designed for single-site setups (Section 5).

2 RELATED WORK

MapReduce applications have been studied extensively on single-cloud computing platforms [7], [8]. Storage elasticity [9], [10] is a particularly interesting aspect for iterative applications, as it is an important component of the overall pay-as-you-go cost. Data shuffling is another difficult problem even in a single data center [11].

Iterative MapReduce applications have been emphasized in [12], [13]. The proposed optimizations included making the task scheduler loop-aware and adding various

caching mechanisms. Although designed for single-site deployment, such techniques can be adapted for hybrid cloud bursting to complement our own techniques. In this context, the core ideas of the performance model proposed in this paper are generic enough to enable an extension that takes such additional optimizations into consideration.

Performance and cost prediction in a single data center have been studied from multiple angles. A storage performance and cost model for iterative applications was introduced in [14]. Given the data-intensive nature of MapReduce applications and the need to persist data between jobs, such a direction is an important complement to our own work. MapReduce performance modeling in particular has focused on various aspects: scheduling, resource provisioning, performance, and cost estimation.

Tian and Chen [15] proposed a cost model that estimates the performance of a job from a set of test runs on a small input datasets and small number of nodes. The authors provisioned the resources for the job using a simple regression technique. Chen et al. [16] further improved the cost model and proposed CRESP, which employs a brute-force search technique for provisioning the optimal cluster resources in term of map slots and reduce slots for Hadoop jobs. In the two models, however, the number of reduce tasks have to be equal to the number of reduce slots; hence, these models consider only a single wave of the reduce phase.

Many tuning frameworks for Hadoop MapReduce have been proposed. Lama and Zhou [17] proposed AROMA, a machine learning optimization framework that automatically selects the cloud resources and tunes the configuration parameters of a Hadoop MapReduce job to achieve quality of service (QoS) goals with as little cost as possible. Similarly, Starfish [18] relies on dynamic Java instrumentation to collect profile information about executed jobs at a fine granularity for job estimation and automatic optimization. Such approaches complement our proposal in that they can act as a tool to find the optimal application-level parameters needed to obtain the performance prediction model.

ARIA [19] and its extension [20] introduce a MapReduce analytical performance model based on the makespan theorem for homogeneous compute clusters. Using this approach, one can reason about performance prediction in terms of upper and lower bounds. Herodotou [21] refined a more detailed mathematical model of each phase of MapReduce, where the map phase is divided into five subphases and the reduce phase is divided into four subphases. As a follow-up, the work presented in [22] estimated the duration of these phases by using regression techniques. The estimated values were then used in the analytical model presented in [19] to estimate the overall job execution time. The same bound-based approach [19] was applied by Zhang et al. [23] to heterogeneous Hadoop cluster environments. Different from our work, heterogeneity is understood as compute nodes with different processing power located in the same cluster (which raises an entirely different set of problems other than network partitioning due to a weak link). Several other works show interest [24]–[26] in heterogeneous MapReduce environments.

Because of the specific nature of hybrid clouds that introduce a weak link and iterative MapReduce applications that reuse input data, none of the previous models is accu-

rate enough or easy to adapt in this context. Therefore, to our best knowledge, we are the first to reveal and solve the challenges associated with performance modeling when both *iterative MapReduce applications* and *hybrid cloud bursting* are considered simultaneously.

3 DATA LOCALITY CHALLENGES AND SOLUTIONS FOR HYBRID CLOUD BURSTING

The MapReduce paradigm is specifically designed to facilitate a high degree of data parallelism. In the first stage (*map phase*), massive amounts of input data are read from a storage layer (typically a distributed file system such as HDFS [27]) and transformed in an embarrassingly parallel fashion by mapper processes such that an intermediate output (consisting of key-value pairs) is obtained that is sorted by key. Then, in a second stage (*shuffle phase*), a separate set of reducer processes fetches (in parallel) the data corresponding to individual keys from all mappers and merges it. In a third stage (*reduce phase*), the reducers apply an aggregation over the values merged under the same key to obtain the final result (one value per key), which is typically persisted in the storage layer.

Both the mappers and reducers are distributed processes that exhibit highly concurrent I/O-intensive data access patterns, which can overwhelm the networking infrastructure with internode data transfers. In order to address this issue, *data locality* awareness is a key feature of MapReduce. The storage layer is colocated with the runtime on the same nodes and is designed to expose the location of the data blocks, effectively enabling the scheduler to bring the computation close to the data and to avoid a majority of the storage-related network traffic.

3.1 Impact of the Weak Link

A hybrid cloud-bursting scenario adds a new level of complexity to data locality, as illustrated in Figure 1a. In this context, the on-premise VMs and off-premise VMs compete for the same wide area link when cross-site communication is necessary. In addition, this link is often of lower throughput and higher latency than the direct inter-VM links. As a consequence, highly concurrent cross-site data transfers are likely to saturate the wide area link and cause an I/O bottleneck. For this reason, we refer to the wide area link as the *weak link*. Specifically, the weak link affects MapReduce applications in the following fashion.

Map phase: Since the input data is present initially only on the on-premise VMs, any map task that is scheduled off-premise needs to access the on-premise data, thus involving a data transfer over the weak link. Furthermore, all off-premise mappers are running in parallel and thus compete for the weak link, thereby introducing high I/O pressure on it.

Shuffle phase: Each reduce task needs to collect the intermediate data generated by the map tasks: if r reduce tasks collect the intermediate data from m map tasks, an m -to- r concurrent communication is required during this phase. Therefore, the weak link will be stressed by all communication required between the on-premise mappers and the off-premise reducers, and between the off-premise mappers and the on-premise reducers, as can be seen in Figure 1b.

Reduce phase: Once the reduce tasks have finished pulling the intermediate data and have performed the aggregation, they typically need to persist the results on-premise. Again, this process involves data transfers from the off-premise VMs to the on-premise VMs over the weak link, which put I/O pressure on it.


In the context of iterative MapReduce applications, the impact of the weak link accumulates as each iteration needs to go through all three phases. However, the iterations are not independent MapReduce jobs since they share a large part of the initial input data. Thus, one can leverage this particular aspect in order to reduce the pressure on the weak link. To this end, we introduced in our previous work [3], [4] two complementary techniques to improve data locality for hybrid cloud bursting. For completeness, we present these techniques below.

3.2 Off-Premise Replication Using Site-Aware Rebalancing

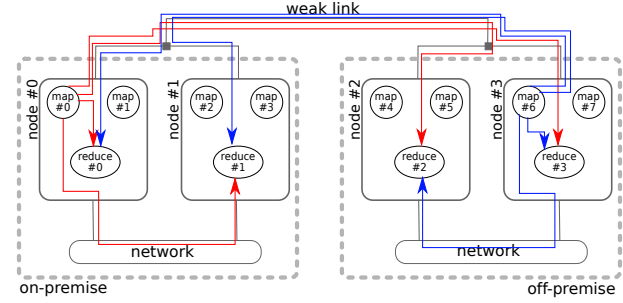
A naive approach that is oblivious to data locality is to simply leave the input data on-premise and pull it on-demand from the off-premise map tasks. While straightforward, this approach has a major disadvantage: the input data needs to be transferred over the weak link unnecessarily for each iteration, leading to performance degradation. A better choice is to ship the input data off-premise before running the iterative MapReduce application. Such an approach avoids putting I/O pressure on the weak link during the runtime of the application. But the initial data transfer is a time-consuming process that leads to extra runtime overhead, while generating extra cost due to additional off-premise storage space utilization.

Therefore, a much better choice is to ship the input data to the off-premise VMs *asynchronously* during runtime. Such an approach hides the data transfer overhead in the background, but it may interfere with the application communication patterns (the VMs act as both data store and compute elements that share the same communication channels). We have shown that input data can be efficiently shipped off-premise by using *rack awareness*. This was initially implemented as a resilience mechanism in HDFS [27], the default storage layer of Hadoop. Specifically, data blocks are replicated in HDFS (three times by default) for resilience purposes, with at least one replica in a different rack from the one where the write originated. Furthermore, HDFS also can rebalance the replicas across the storage elements to distribute the load evenly while preserving rack awareness. Thus, by deploying new HDFS storage elements on the off-premise VMs as a separate rack, a rebalancing operation will migrate one replica for each data block to the off-premise VMs asynchronously, effectively achieving a *site-aware rebalancing*.

The main advantage of this approach is that it minimizes the amount of data transferred off-premise (a single replica) to achieve full potential of exploiting the data locality, while maintaining the resilience properties. Furthermore, although illustrated on HDFS, the core idea (asynchronous migration of a single replica) can be implemented also for in-memory frameworks (e.g., Spark [28], Pregel [29]) that aim to minimize the interaction with the storage layer.



(a) Infrastructure: The weak link interconnects the on-premise VMs with the off-premise VMs



(b) MapReduce shuffle phase: Concurrent data transfers put I/O pressure on the weak link. For clarity, only data transfers from one on-premise map and one off-premise map are depicted.

Fig. 1: Schematic representation of a hybrid cloud-bursting architecture and its implications for MapReduce applications

3.3 Scheduling Based on Enforced Multilevel Locality

The default Hadoop scheduler uses data locality only as a preferential matching mechanism between map tasks and free slots. If asynchronous HDFS rebalancing is employed, however, a map task may be scheduled off-premise before a replica of its corresponding data block was migrated, thus triggering a pull and leading to a double data transfer of the same block. In this case, delaying the scheduling of such off-premise map tasks is beneficial, under the assumption that avoiding stress on the weak link leads to a smaller overall overhead. To this end, we propose an enforced multilevel locality scheduling policy. Specifically, a map task will be preferentially scheduled on the same VM that holds its data dependencies. If all VMs that satisfy its data dependencies are busy, then the map task will be scheduled to another off-premise VM with a free compute slot only if the data was already migrated off-premise. Otherwise, it will be scheduled on an on-premise VM even if all on-premise VMs are busy. With this approach, off-premise map tasks will never pull data remotely over the weak link, and thus interferences with the asynchronous background data migration will be avoided.

We implemented this policy in Hadoop by modifying the resource manager to make use of the *relaxLocality* flag. As in the previous case, the core idea of enforced multilocality can be extended beyond Hadoop MapReduce to in-memory processing frameworks as well.

4 PERFORMANCE MODELING PROPOSAL

In this section, we introduce the broad principles and methodology behind our performance modeling proposal. Our goal is to estimate the completion time of a given iterative MapReduce application for hybrid cloud bursting based on a series of system-level and application-level parameters that are extracted in advance.

Given the locality challenges of hybrid cloud bursting, our proposal assumes that the underlying big data analytics framework for which the user needs to estimate the completion time is optimized by using the principles discussed in Section 3. Specifically, we make two assumptions: (1) there is a data rebalancing that migrates one replica of the input data asynchronously off-premise while the application is running; and (2) the scheduling policy for the

map tasks obeys the enforced multilevel locality condition. These assumptions influence the mathematical expressions we introduce in Section 4.4.

For simplification, we assume that the setup consists of a fixed set of on-premise VMs and off-premise VMs of similar capability, which gives us a fixed set of system-level parameters. Users interested in estimating the runtime for various setups (e.g., finding the optimal number of off-premise VMs to achieve the desired runtime) can apply our approach for each configuration individually.

Our approach consists of three steps. First, we run a synthetic benchmark to extract the fixed system-level parameters corresponding to the on-premise and off-premise VMs. These parameters are *independent of the application* and can be reused for a different application or user (e.g., they can be cached on-premise). We refer to this step as *calibration*.

Second, we extract the application-level parameters. These parameters are *independent of the hybrid setup* and can be either known in advance or obtained by running the application at smaller scale on-premise only. In this way, users can estimate the benefits of hybrid cloud bursting without actually ever trying it, as long as the calibration step was already performed for the desired configuration. We refer to this step as *characterization*.

Third, once both the calibration and the characterization is complete, we apply a mathematical expression to estimate the completion time. Note that the approximations introduced above have an inherent variability due to the complexity of MapReduce applications in general and the additional complexity introduced by the weak link. Therefore, it is important to be able to present both the optimistic (*lower bound*) and pessimistic (*upper bound*) runtime estimation to the user. To derive the mathematical expressions for both cases, we use the makespan theory as applied in the context of MapReduce.

To aid the extraction of both system-level and application-level parameters, we have developed a tool that can analyze a job in terms of map/shuffle/reduce times; HDFS data distribution; task distribution between on-premise and off-premise; and node statistics such as CPU, I/O network, I/O disk, and memory utilization. Our tool extracts information from a combination of Hadoop counters, Hadoop logs, Hadoop Rumen tool [30], and Systat [31]

to generate profiling information.

Note that while we illustrated the core idea of our proposal in the context of Hadoop MapReduce, its applicability can be extended to other big data analytics frameworks that are based on the same fundamentals [28], [29], [32] (colocate data storage with computations, split a computations into subtasks, schedule subtasks close to the data). Also, the generic aspect of the additional assumptions specific to hybrid cloud bursting that our proposal is based on were explained in Section 3.

In the rest of this section, we detail the three-step proposal summarized above.

4.1 Theoretical Makespan Bounds for MapReduce

For completeness, we briefly introduce the theoretical makespan bounds as applied in the context of MapReduce. This presentation provides preliminary background needed in order to understand the reasoning behind the parameters extracted from the calibration and characterization steps.

Each of the map and reduce phases can be abstracted as a series of n tasks of duration t_i (where i refers to each task) that need to be processed by k slots. The assignment of tasks to slots is done dynamically by the MapReduce runtime according to a simple greedy policy that assigns each task to the slot with the earliest finishing time.

The best case is obtained when all slots are evenly loaded, in which case each slot is busy for at least $(\sum_{i=1}^n t_i)/k = (t_{avg} \times n)/k$. Therefore, this is the lower bound of the makespan.

The worst case is obtained when the longest task t_{max} is scheduled last. This means that the k slots are busy with the other tasks and take at most $(t_{avg} \times (n - 1))/k$ to process them. Once they finish with these tasks, one of them needs to process t_{max} . Therefore, the upper bound of the makespan is $(t_{avg} \times (n - 1))/k + t_{max}$.

We note that the lower bound can be expressed more precisely as a function of the average task duration (rather than the minimum). Consequently, several of the parameters we extract from the calibration and characterization steps are averages and maximum values.

4.2 Calibration Using Synthetic Benchmarking

We develop a synthetic benchmark that focuses on the extraction of the I/O and communication overheads in a hybrid setup. Specifically, the goal is to extract these overheads for each phase (map, shuffle, reduce) based on the quantity of data involved: (1) the amount of data read from HDFS; (2) the amount of data written to HDFS; and (3) the amount of network traffic between map and reduce tasks. Using this approach, we can cover all possible combinations of data sizes in each phase.

To achieve this goal, we implement a collection of map, combiner, and reduce functions that generate a synthetic workload based on a series of configurable input parameters used to specify the amounts of data. Both the map phase and the reduce phase of the synthetic workload deliberately avoid computational overhead (minimal load on the CPU) in order to isolate the I/O and communication overheads. The map phase is structured in two parts: it reads the input chunk in the map function and writes a specified amount of

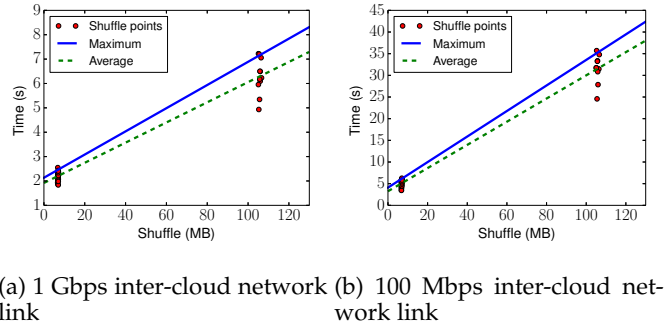


Fig. 2: Shuffle approximation in a 3-ON-3-OFF premise hybrid architecture




Fig. 3: Reduce approximation in a 3-ON-3-OFF premise hybrid architecture

intermediate data in the combiner function. The output of the combiner is grouped by key, with each group collected by the corresponding reducer. The reducer simply writes a predefined amount of data as output in HDFS.

This synthetic benchmark is then executed for a variable quantity of data in a hybrid setup comprising a fixed number of on-premise and off-premise VMs. To illustrate how this *calibration* process works, we present an experimental example based on a typical hybrid cloud setup (described in more detail in Section 5.1), where we use two representative configurations for the weak link: 100 Mbps and 1 Gbps.

Then, based on the profiling information, we define a series of system-level parameters that quantify the hybrid-specific overheads.

Hybrid map overhead stretch coefficient: Although the scheduler forces the execution of mappers on the nodes where the input data from HDFS is present, there are some extra reads when the logical end of data does not exactly match the HDFS end of a split (i.e., the input of the mapper is covered by two different HDFS chunks). Therefore, mappers will sometimes read remote data over the weak link, creating extra overhead compared with the on-premise-only case. We measure this overhead using the synthetic benchmark and express it as a stretch coefficient denoted α .

Hybrid approximation of shuffle overhead per reducer: Once a sufficient number of mappers have finished producing the intermediate data, reduce tasks are launched and begin collecting the data. However, a reduce task cannot start the aggregation at the same time as the intermediate

data is collected because it needs to sort the data first (which cannot happen before all mappers have finished and their intermediate data was collected). Therefore, each reducer experiences a shuffle overhead that is proportional to the amount shuffle data it needs to pull. Since the amount of shuffle data per reducer (denoted d_{Sh}) is application dependent, we express the shuffle overhead as an approximation function (denoted f_{Sh}). To this end, we choose a set of representative shuffle sizes, measure the shuffle overhead per reducer using the synthetic benchmark, and then apply linear regression to obtain two approximation functions: one for the average (using all reducers) and one for the maximum (using only the slowest reducer for each shuffle size). Note that we have chosen linear regression because it was confirmed by previous work as a good approximation for shuffle behavior [22], [33], [34] for a variety of network topologies with different performance characteristics. In our case, we assume that the weak link will slow data transfers uniformly, thus preserving the linear behavior. Nevertheless, our proposal is general enough to work with any approximation function other than linear regression, should the need arise (e.g., limited weak link or low QoS with performance fluctuations). This aspect is also subject to a trade-off: the more complex the approximation function and the number of points needed to obtain it, the longer the calibration will take.

Figure 2 illustrates this for a hybrid scenario with three on-premise VMs and three off-premise VMs using 24 mappers and 12 reducer slots. The actual runtimes of the reducer tasks are illustrated as clusters of points, while the approximation function (average and maximum) is illustrated as a line.

Hybrid approximation of write overhead per reducer: In the reduce phase, the computational overhead of each reducer does not depend on the weak link and represents a significant part of the runtime of the reduce phase. Once the reducer has finished the computation, however, it needs to write the output to the storage layer, thereby stressing the weak link. Again, the amount of output per reducer (denoted d_R) is application dependent. Thus, we need to express it as an approximation function (denoted f_{Rd}) in a manner similar to the shuffle overhead. Using linear regression, we obtain the corresponding approximation function. Figure 3 illustrates this for the same hybrid scenario used above (both average and maximum). Again, linear regression can be replaced with another approximation function if necessary.

Rebalancing bandwidth: In addition to the overheads related strictly to the MapReduce runtime, estimating how long the off-premise rebalancing of the input data will last is important because the process runs asynchronously in the background while the iterative MapReduce application is progressing, thereby creating interference and potentially visible slowdown. To this end, we run a HDFS rebalancing using a large HDFS input data size (e.g., 10 GB) and measure its completion time. Then, we compute an average bandwidth, denoted β , that is used in our performance model to account for the effect of the rebalancing.

Table 1 summarizes all system-level parameters extracted by using the calibration step.

TABLE 1: Parameters obtained from the calibration and characterization.

System-level parameters		
Name	Description	
S_{on}^M	On-premise map slots	
S_{off}^M	Off-premise map slots	
S^R	Reduce slots	
α	Hybrid map phase stretch factor	
β	Rebalancing bandwidth	
$f_{Sh}(d_{Sh})$	$f_{Sh}^{max}(d_{Sh})$	Shuffle time per reducer (avg and max)
$f_{Rd}(d_{Rd})$	$f_{Rd}^{max}(d_{Rd})$	Write time per reducer (avg and max)
Application-level parameters		
Name	Description	
M	Total number of mappers	
R	Total number of reducers	
t_{Mp} and t_{Mp}^{max}	Total runtime for mappers (avg and max)	
t_{Rc} and t_{Rc}^{max}	Sort and aggr. for reducers (avg and max)	
d_{Sh}	Amount of shuffle data per reducer	
d_{Rd}	Amount of output data per reducer	

4.3 Application Characterization

In this section we show how to extract the necessary information to characterize the iterative MapReduce application. Note that this step is necessary only when the application-level parameters introduced below (needed by the mathematical expression) are not known in advance or cannot be directly computed based on some existing knowledge about the application.

Specifically, the user needs to run the actual application on-premise only at reduced scale (both number of nodes and number of iterations) and extract the following information

- Total number of mappers: M
- Total number of reducers: R
- Average and maximum time to run a mapper (refers to the total time that includes both CPU and I/O): t_{Mp} and, respectively, t_{Mp}^{max}
- Average and maximum time to run a reducer computation (includes the time to sort and compute the aggregation, but not the time to write the result to HDFS): t_{Rc} and, respectively, t_{Rc}^{max}
- Amount of input data in the shuffle phase, d_{Sh}
- Amount of output data of the reduce phase, d_{Rd}

Note that the MapReduce framework incurs a scheduling overhead, which is observable as a gap between the tasks assigned to the mapper and reducer slots. For simplification, we assume that t_{Mp} and t_{Rc} already include the average gap duration, while t_{Mp}^{max} and t_{Rc}^{max} already include the maximum gap duration.

Table 1 summarizes these application-level parameters.

4.4 Performance Model

In this section, we introduce a performance model that enables users to estimate the runtime of iterative MapReduce applications in hybrid cloud-bursting scenarios. The performance model is a mathematical expression that uses the application-level and system-level parameters described in Table 1 as input and produces an estimation of the total runtime as output.

The lower bound of the completion time $Total$ of an iterative map-reduce job with n iterations, can be expressed as

$$Total = \sum_{i=1}^n T(i). \quad (1)$$

$T(i)$ is the lower bound of the i th iteration and can be decomposed as

$$T(i) = T_{Mp}(i) + T_{Sh}(i) + T_{Rd}(i), \quad (2)$$

where T_{Mp} is the lower bound for the map phase, T_{Sh} for the shuffle phase, and T_{Rd} for the reduce phase.

Similarly, for the upper bound we obtain

$$Total^{max} = \sum_{i=1}^n T^{max}(i) \quad (3)$$

$$T^{max}(i) = T_{Mp}^{max}(i) + T_{Sh}^{max}(i) + T_{Rd}^{max}(i). \quad (4)$$

For the rest of this section, we detail how to obtain each of T_{Mp} , T_{Sh} , and T_{Rd} .

4.4.1 Completion time of map phase

To obtain a mathematical expression to estimate the completion time of the map phase, one must understand how this phase evolves during the successive iterations of the MapReduce job. We force the execution of any map task on a node where there is a copy of its input data, so for the first iteration all the map tasks will be scheduled on-premise only. In parallel with the execution of the first iteration, the rebalancing of the input data to the off-premise infrastructure proceeds in the background. Thus, at the beginning of the second iteration, some replicas of the input chunks have already been migrated on the off-premise nodes, and the scheduler will launch off-premise mappers to handle them. As the rebalancing progresses, the number of map tasks that will be executed off-premise will increase at each iteration until it stabilizes (which is not necessarily the moment when the rebalancing has finished, because the off-premise map slots can be saturated even before a replica of each chunk was transferred off-premise).

We express this intuition mathematically as follows: for the first iteration, all map tasks (M) will be executed on-premise. To simplify our equations, we use t_{Mp} and t_{Mp}^{max} as a shortcut to denote $t_{Mp} \times \alpha$ and $t_{Mp}^{max} \times \alpha$, respectively. Since there are S_{on}^M map slots on the on-premise nodes, the lower bound of the map phase for the first iteration is

$$T_{Mp}(1) = \frac{M}{S_{on}^M} \times t_{Mp}. \quad (5)$$

For the second iteration, M_1^{off} input chunks will be present off-premise, having been transferred during the previous iteration. Therefore, M_1^{off} map tasks can be executed off-premise. This number can be approximated by using the rebalancing bandwidth (β), the size of the HDFS chunk s and the runtime of the first iteration $T(1)$ as follows:

$$M_1^{off} = \frac{\beta \times T(1)}{s}. \quad (6)$$

If M_1^{off} map tasks are scheduled off-premise, the rest of map tasks, $M - M_1^{off}$ (which we assume is greater

than M_1^{off}), will be scheduled on-premise. Thus, the lower bound of the map phase runtime of the second iteration is

$$T_{Mp}(2) = \frac{M - M_1^{off}}{S_{on}^M} \times t_{Mp}. \quad (7)$$

Using the same reasoning for the third iteration, we obtain for the lower bound of the map phase

$$T_{Mp}(3) = \frac{M - ((\beta \times (T(2) + T(1)))/s)}{S_{on}^M} \times t_{Mp}. \quad (8)$$

By generalization, for the i iteration we obtain the following:

$$T_{Mp}(i) = \frac{M - ((\beta \times \sum_{j=1}^{i-1} T(j))/s)}{S_{on}^M} \times t_{Mp} \quad (9)$$

This formula is true as long as all on-premise map slots are filled and there are off-premise idle map slots that cannot be used because the rebalancing did not ship enough chunk replicas off-premise. The moment when stabilization happens can be expressed mathematically as follows:

$$M^{off} \geq \frac{S_{off}^M}{(S_{on}^M + S_{off}^M)} \times M. \quad (10)$$

From this moment on, the time to process on-premise scheduled maps is almost the same as the time to process off-premise maps. In this situation, the number of map tasks scheduled off-premise will be M^{off} and the number of map tasks scheduled on-premise $M - M^{off}$. These numbers will remain constant for the rest of the iterations, leading to the following expressions for the remaining runtime:

$$T_{Mp}(i) = \frac{M - M^{off}}{S_{on}^M} \times t_{Mp} \text{ where} \quad (11)$$

$$M^{off} = \frac{S_{off}^M}{(S_{on}^M + S_{off}^M)} \times M.$$

The expressions in Equations 9 and 11 can be joined into a single expression that estimates the lower bound of the map phase for any iteration:

$$T_{Mp}(i) = \frac{M - M_{i-1}^{off}}{S_{on}^M} \times t_{Mp} \text{ where} \quad (12)$$

$$M_{i-1}^{off} = \min \left(\frac{\beta \times \sum_{j=1}^{i-1} T(j)}{s}, \frac{S_{off}^M}{S_{on}^M + S_{off}^M} \times M \right).$$

Upper bound for the map phase: By applying the theoretical makespan results detailed in Section 4.1 to the lower bound described in Equation 12, we obtain the following expression for the upper bound:

$$T_{Mp}^{max}(i) = \left(\left\lceil \frac{M - M_{i-1}^{off}}{S_{on}^M} \right\rceil - 1 \right) \times t_{Mp} + t_{Mp}^{max}. \quad (13)$$

4.4.2 Completion time of the shuffle phase

The shuffle phase is entirely managed by the MapReduce runtime and has no application-specific overhead. Therefore, to estimate the lower bound of the shuffle phase we simply need to apply the shuffle phase approximation function f_{Sh} (obtained from the calibration step) to the shuffle size per reducer d_{Sh} (obtained from the application characterization). We obtain the following:

$$T_{Sh}(i) = \frac{R}{SR} \times f_{Sh}(d_{Sh}), \forall i = 1..n. \quad (14)$$

Similarly, for the upper bound we apply the makespan results to obtain

$$T_{Sh}^{max}(i) = \left(\left\lceil \frac{R}{SR} \right\rceil - 1 \right) \times f_{Sh}(d_{Sh}) + f_{Sh}^{max}(d_{Sh}), \forall i = 1..n. \quad (15)$$

4.4.3 Completion time of the reduce phase

The reduce phase consists of a number of reducers R that compete for a number of parallel reducer slots S^R . In this case, the average completion time of a reducer t_{Rd} depends on both the application-level parameters and the system-level parameters. Specifically,

$$t_{Rd} = t_{Rc} + f_{Rd}(d_{Rd}). \quad (16)$$

The explanation is the following. There is an application-dependent computational (t_{Rc}) overhead (obtained through characterization), in addition to the HDFS write overhead (obtained by applying the application agnostic approximation function f_{Rd} to the application-specific output size per reducer d_{Rd}).

Thus, the lower bound of the reduce phase can be estimated as follows:

$$T_{Rd}(i) = \frac{R}{SR} \times t_{Rd}, \forall i = 1..n. \quad (17)$$

Similarly, for the maximum, the following applies:

$$T_{Rd}^{max} = t_{Rd}^{max} + f_{Rd}^{max}(d_{Rd}). \quad (18)$$

By applying the makespan results, we obtain the following upper bound:

$$T_{Rd}^{max}(i) = \left(\left\lceil \frac{R}{SR} \right\rceil - 1 \right) \times t_{Rd} + t_{Rd}^{max}, \forall i = 1..n. \quad (19)$$

4.5 Complex Iterations

So far we have made an important assumption about the applications: each iteration involves a single MapReduce job that is computationally similar to the previous iterations. In practice, however, iterations can be complex and involve a series of steps expressed as separate MapReduce jobs (e.g., PageRank, as described in Section 5.4).

In this section we briefly show how to generalize our approach to address such complex iterations. Let m be the number of MapReduce jobs in a complex iteration i and $T_j(i)$ the runtime of the j -th MapReduce job in the sequence of m jobs. Then, the runtime of each complex iteration i is the sum of the durations of the m MapReduce jobs:


$$T(i) = \sum_{j=1}^m T_j(i). \quad (20)$$

For n iterations, the following holds:


$$Total = \sum_{i=1}^n \left(\sum_{j=1}^m T_j(i) \right) = \sum_{j=1}^m \left(\sum_{i=1}^n T_j(i) \right) = \sum_{j=1}^m Total_j. \quad (21)$$

In other words, we can see an application with complex iterations as the equivalent serialization of m sub-applications with simple iterations. In this case, we can simply characterize each of the m sub-applications, apply our mathematical expressions to estimate their individual completion time, and sum up the estimations to obtain the final estimation for the original application with complex iterations. Since the system-level parameters are application agnostic, the calibration needs to be performed only once regardless of m .

As an example, we illustrate this process for an application with four complex iterations, each of which is composed of three MapReduce jobs.



Then, the equivalent serialization yields three sub-applications with four simple iterations (single-MapReduce jobs), for which our approach can be individually applied.



5 EVALUATION

In this section we evaluate the effectiveness of our approach experimentally, using a variety of scenarios and comparisons that involve multiple real-life iterative MapReduce applications.

5.1 Experimental Setup

The experiments for this work were performed on the *Kinton* testbed of the HPC&A group based at Universitat Jaume I. It consists of eight nodes, all of which are interconnected with 1 Gbps network links and split into two groups: four nodes feature an Intel Xeon X3430 CPU (4 cores), HDD local storage of 500 GB, and 4 GB of RAM. These less powerful nodes (henceforth called *thin*) are used for management tasks. The other four nodes feature two Intel Xeon E5-2630v3 (2 x 8 Cores), HDD local storage of 1 TB, and 64 GB of RAM. These more powerful nodes (henceforth called *fat*) are used to host the VMs.

We configure two separate IaaS clouds: *on-premise* and *off-premise*, each running its separate OpenStack Icehouse instance (QEMU/KVM as the hypervisor). The VM instances of the same cloud are configured to directly communicate with each other via the links of their compute node hosts. However, all communication outside of the same cloud is passing through a dedicated network node (Neutron) that acts as a proxy and is part of the default OpenStack distribution. Thus, in a hybrid OpenStack setup the weak

link is defined by the end-to-end bandwidth between the two proxies. We control the available bandwidth to cover two representative settings: 1 Gbps and 100 Mbps. These correspond to the case when the user decides to buy premium access to the cloud (i.e., dedicated fast link) vs. regular access.

5.2 Approaches

We compare four approaches throughout our evaluation.

On-Premise Actual: Corresponds to the case when all VMs are on-premise and no weak link can cause an I/O bottleneck. In this case, a standard Hadoop deployment is used. We use it as a lower bound for comparison, showcasing what would happen in an ideal scenario where the user has no cost constraints and can afford to invest in additional on-premise resources to achieve the highest performance rather than adopt a hybrid solution.

ARIA (Automatic Resource Inference and Allocation): Is a state-of-the-art framework that estimates the runtime of a single-site MapReduce application based on its profile (application-level parameters) and then optimally schedules it to meet a given soft deadline [19]. To our best knowledge, no performance prediction approach specific for hybrid cloud bursting is available, so ARIA would be the closest estimation users can achieve with state of the art. Therefore, a comparison with ARIA is highly relevant to show that single-site techniques for estimating the runtime of iterative MapReduce applications are not accurate enough for use in a hybrid cloud-bursting scenario, thus justifying the need for specialized models like our proposal.

Hybrid Actual Runtime: Corresponds to the real measured runtime of an iterative MapReduce application using a given on-premise and off-premise configuration of VMs. The Hadoop deployment used to run application is optimized for hybrid cloud bursting using the rack-local scheduling and asynchronous rebalance techniques described in Section 3.1. We use this approach for comparison in order to showcase the accuracy of the estimations provided by our approach.

Hybrid Estimated Runtime: Corresponds to the estimated runtime of an iterative MapReduce job using our proposal (Section 4), which is optimized for a hybrid cloud-bursting scenario where the Hadoop deployment employs rack-local scheduling and asynchronous rebalance (Section 3.1).

5.3 Methodology

For our experiments, we created a new VM flavor with 4 vCPUs, HDD local storage of 100 GB, and 16 GB of RAM. Thus, each compute node has the capacity to host 4 VMs simultaneously. Since some VMs are colocated on the same node, the virtual network interface of all VMs is limited to 1 Gbps, in order to avoid differences between VMs hosted on the same node vs. remote nodes. We use one fat node to provision up to 4 VMs on the on-premise part and three nodes to provision up to 12 VMs on the off-premise part. We deploy Hadoop 2.6.0 initially on-premise only: one VM is used as the Hadoop master (both MapReduce and HDFS), the rest of the VMs are used as Hadoop slaves (both MapReduce and HDFS). Each Hadoop slave is configured

with enough capacity to run four mappers and four reducers simultaneously. Any initial input data is stored on-premise only in the initial HDFS deployment.

First, we run the application on-premise only and record the runtime, for the smallest case (3 VMs). We call this the *baseline* case. An important premise for any other setup (whether on-premise or hybrid) is to show speed-up with respect to the baseline (otherwise it does not make sense to commit more VMs).

Then, using the profiling information, we extract the application-level parameters described in the characterization step (Section 4.3). For completeness, we also run *strong scalability* experiment (constant problem size) by increasing number of VMs from 6 up to 15 in steps of 3. This experiment is not involved in the extraction of the application-level parameters but facilitates the study of the results and corresponds to the *on-premise actual* case.

Second, we use the application-level parameters in order to estimate the runtime with ARIA. Again, we apply ARIA for an increasing number of VMs to show how the runtime scales in comparison with the baseline.

Third, we run another strong-scalability experiment where we keep the number of on-premise VMs fixed at three, while adding an increasing number of off-premise VMs: from 3 up to 12 in steps of 3. For each resulting hybrid bursting scenario, we run (1) the calibration (using the generic benchmark introduced in Section 4.2) to extract the system-level parameters and (2) each application to obtain the *hybrid actual runtime*.

We then use both the application-level and system-level parameters to estimate the runtime using our proposal for each hybrid cloud-bursting scenario, which yields the *hybrid estimated runtime*. Note that many applications we study exhibit complex iterations composed of multiple MapReduce jobs. In this case, we apply the observations from Section 4.5 to compute the hybrid estimated runtime.

5.4 Applications

For this work, we use four representative iterative MapReduce applications selected from industry-standard big data benchmarks. We selected these applications to cover as much diversity as possible, in terms of both application domain and MapReduce behavior (map-intensive, reduce-intensive, or both).

Iterative Grep (I-GREP): Is a popular analytics tool for large unstructured text, where at each iteration a new set of keywords is used to narrow the search in relevant text obtained from previous iterations. For example, one may want to count how many times a certain concept is present in the Wikipedia articles and, depending on the result, prepare the next regular expression in order to find correlations with another concept. Since the regular expression is typically an exact pattern, the output of the mappers simply consists of a small number of key-value pairs that are reduced to a single key-value pair. Thus, it can be classified as a typical map-intensive job. The implementation is based on `grep`, which is included with the Hadoop distribution.

KMeans: Is a widely used application for vector quantization in signal processing, cluster analysis in data mining, pattern classification, and feature extraction for

machine learning [35]. It is based on iterative refinement: each iteration aims to improve the partitioning of a multidimensional vector into k clusters such that sum of squares of distances between all vectors of the same cluster and their mean is minimized. This process repeats until the improvement obtained during an iteration is smaller than a predefined epsilon. K-Means was shown to be efficiently parallelizable and scales well using MapReduce [36]. From a data-management perspective, it is a good example of a map-intensive application that reuses the initial input data at each iteration. The implementation we use is part of the Intel’s *HiBench* [37] big data benchmark.

PageRank: Is a link analysis algorithm [38] that assigns a numerical weight to each element of a hyperlinked set of documents, (e.g., WWW) with the purpose of quantifying its relative importance within the set. It is widely used in web search engines to calculate the ranks of web pages in function of the number of reference links. Its iterative nature is more complex and involves two successive MapReduce jobs: (1) an output-intensive phase where the reduce phase generates twice as much data as the input data read by the map phase and (2) a shuffle-intensive phase where the output of the mappers is equal in size to the input. Thus, PageRank is a good example of a balanced application with complex iterations that is both map intensive and reduce intensive, while generating a lot of intermediate data that is not reused. Again, the implementation we use is part of the Intel’s *HiBench* big data benchmark.

Connected Components: Is a well-known graph problem arising in a large number of applications including data mining, analysis of social networks, image analysis, and related problems. It aims at identifying groups of connected vertices in a graph by iteratively assigning each vertex to its smallest neighbor to form starlike subgraphs [39]. Connected Components is a good example of a reduce-intensive application. The implementation we use is part of Intel’s *BigBench* [40] benchmark, which also includes a data generator.

5.5 Results

Using the methodology presented in Section 5.3, we perform an experimental study for each of the real-world applications described in Section 5.4. For all runtime estimations, we compute both the lower and upper bound and derive the average from these bounds. In addition, we study the accuracy of all average estimations (vs. the actual runtime) at fine grain for both the map phase and the reduce phase separately, which provides additional insight into the overall accuracy. For the rest of this section, we discuss the results for each application individually.

The first application we study is I-GREP. We use as input data 20 GB worth of Wikipedia articles, which are queried successively in 50 iterations using 50 different keywords. Each iteration is complex and is composed of two jobs per iteration (search and sort stages). The shuffle data for each iteration is less than 1 MB, which means the map phase dominates the runtime. The baseline case (3 VMs on-premise) has a runtime of 6483 s. Doubling the amount of on-premise VMs leads to 60% less actual runtime, which shows I-GREP has a good scalability potential.

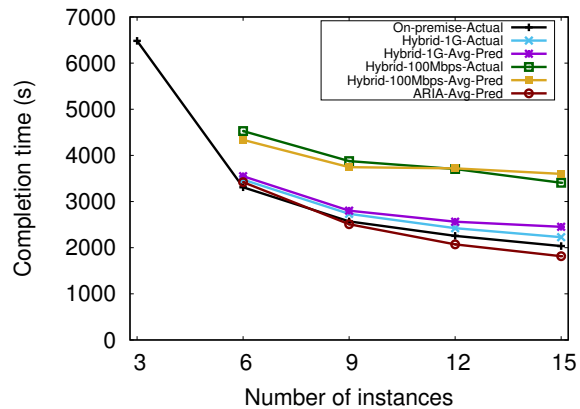


Fig. 4: I-GREP average runtime estimation

Table 2a shows the actual runtimes and the estimations for the upper and lower bounds. Because of high variability between the map tasks, ARIA shows large differences: for 15 VMs it overestimates the upper bound by 30% and underestimates the lower bound by 50% with respect to on-premise actual. Despite this variability, our approach has better accuracy: for 15 VMs and 100 Mbps weak link, it overestimates the upper bound by 20% and underestimates the lower bound by 10%. For 1 Gbps, it overestimates the upper bound by 30% and underestimates the lower bound by 13%.

In Figure 4 we depict the average estimations. As can be observed, despite large difference between the ARIA upper and lower bound, the average estimation is much closer to the on-premise actual values. However, the average ARIA estimation produces large errors against the hybrid scenarios: up to 18.5% for 1 Gbps and up to 46.7% for 100 Mbps. These contrast with the errors produced by our approach: up to 10.1% for 1 Gbps and up to 5.6% for 100 Mbps. We note the increasing accuracy of our approach when the weak link is of low capacity, whereas the opposite holds for ARIA.

Table 2b shows finer-grained accuracy details about the map and reduce phases. As can be observed, in the 100 Mbps weak link case, the ARIA average estimation has a large error for both phases, which explains the overall error. For the 1 Gbps weak link case, the map phase has a small error, but the reduce phase exhibits a large error.

Second, we study the K-Means application. We generate 20 GB of input using the included data generator, which is processed by K-Means in 30 simple iterations. The baseline in this case (3 VMs on-premise) has a runtime of 6471 s. Doubling the number of on-premise VMs leads to 70% less actual runtime, which shows excellent scalability potential.

Table 3a shows the runtime estimations for the upper and lower bound. Unlike the I-GREP case, K-Means exhibits less variability between the map tasks, which improves the accuracy of the ARIA upper and lower bound estimations with respect to actual on-premise. By comparison, our approach has much closer upper and lower bound estimations with respect to hybrid actual.

Figure 5 depicts the average estimations. The ARIA estimation produces a large error for the hybrid 100 Mbps

Approach	6 VMs	9 VMs	12 VMs	15 VMs
On-premise Actual	3310	2569	2256	2035
ARIA Upper Bound	4226	3324	2892	2639
ARIA Lower Bound	2615	1692	1251	992
100 Mbps weak link (3 VMs on-premise)				
Hybrid Actual	4528	3876	3705	3407
Hybrid Upper Bound	4795	4282	4345	4141
Hybrid Lower Bound	3877	3208	3093	3057
1 Gbps weak link (3 VMs on-premise)				
Hybrid Actual	3474	2734	2422	2227
Hybrid Upper Bound	3953	3299	3087	2947
Hybrid Lower Bound	3144	2307	2038	1959

(a) Total actual runtime (s) vs. predicted runtime (s) expressed as upper bound and lower bound

TABLE 2: I-GREP: Map-intensive example of an iterative MapReduce application

Approach	6 VMs	9 VMs	12 VMs	15 VMs
On-premise Actual	3024	2159	1786	1511
ARIA Upper Bound	3501	2499	2020	1739
ARIA Lower Bound	2870	1857	1372	1089
100 Mbps weak link (3 VMs on-premise)				
Hybrid Actual	3872	3207	2893	2743
Hybrid Upper Bound	3900	3374	3266	3091
Hybrid Lower Bound	3649	2997	2799	2713
1 Gbps weak link (3 VMs on-premise)				
Hybrid Actual	3175	2391	1992	1685
Hybrid Upper Bound	3282	2495	2168	1858
Hybrid Lower Bound	3098	2137	1741	1543

(a) Total actual runtime (s) vs. predicted runtime (s) expressed as upper bound and lower bound

Prediction	Accuracy vs. Actual	6 VMs	9 VMs	12 VMs	15 VMs
100 Mbps weak link (3 VMs on-premise)					
ARIA Avg.	(%) Map Error	-18.2	-30.0	-40.0	-43.1
	(%) Red. Error	-56.3	-57.4	-59.4	-59.0
	(%) Total Error	-24.5	-35.3	-44.1	-46.7
Hybrid Avg.	(%) Map Error	-2.6	1.9	5.1	13.4
	(%) Red. Error	-12.4	-25.3	-17.2	-20.8
	(%) Total Error	-4.2	-3.4	0.4	5.6
1 Gbps weak link (3 VMs on-premise)					
ARIA Avg.	(%) Map Error	4.1	-1.8	-8.4	-12.8
	(%) Red. Error	-34.9	-36.5	-37.2	-37.5
	(%) Total Error	-1.5	-8.3	-14.5	-18.5
Hybrid Avg.	(%) Map Error	3.9	2.3	5.6	8.7
	(%) Red. Error	-8.3	3.6	6.7	-15.3
	(%) Total Error	2.1	2.5	5.8	10.1

(b) Accuracy of the average prediction (between lower and upper bound) vs. the hybrid actual runtime broken down by phase

Prediction	Accuracy vs. Actual	6 VMs	9 VMs	12 VMs	15 VMs
100 Mbps weak link (3 VMs on-premise)					
ARIA Avg.	(%) Map Error	-14.7	-29.3	-38.9	-46.9
	(%) Red. Error	-61.6	-63.9	-65.1	-63.9
	(%) Total Error	-17.7	-32.1	-41.4	-48.5
Hybrid Avg.	(%) Map Error	-1.6	1.8	7.5	8.6
	(%) Red. Error	-15.7	-29.1	-20.8	-21.4
	(%) Total Error	-2.5	-0.7	4.8	5.8
1 Gbps weak link (3 VMs on-premise)					
ARIA Avg.	(%) Map Error	2.6	-6.3	-12.0	-12.9
	(%) Red. Error	-41.5	-43.5	-45.1	-45.2
	(%) Total Error	0.3	-8.9	-14.9	-16.1
Hybrid Avg.	(%) Map Error	1.0	-3.5	-2.3	-0.2
	(%) Red. Error	-8.8	1.4	2.6	11.0
	(%) Total Error	0.5	-3.1	-1.9	0.9

(b) Accuracy of the average prediction (between lower and upper bound) vs. the hybrid actual runtime broken down by phase

TABLE 3: KMeans: Map-intensive example of an iterative MapReduce application




Fig. 5: K-Means average runtime estimation




Fig. 6: PageRank average runtime estimation

weak link scenario, reaching almost 50%. For the 1 Gbps case, the error is smaller but still significant at 16%. Our approach reduces the error by an order of magnitude: 5% for the 100 Mbps case and less than 1% for the 1 Gbps case.

Table 3b shows finer-grained details about the accuracy of the map and reduce phases. As can be observed, ARIA has low accuracy in the 100 Mbps case for both phases. In the 1 Gbps case, ARIA has low accuracy for the reduce phase. By contrast, our approach has good accuracy for both phases regardless of the weak link capacity.

Third, we study the PageRank application. We generate 2.8 GB of web data hyperlinks that is processed in 5 complex iterations (2 jobs per iteration). Again, we can see a good scalability potential: the baseline runtime of 3145 s (3 VMs on-premise) is reduced by 51% when doubling the number of on-premise VMs.

The runtime results are listed in Table 4a. In this case, the accuracy of the ARIA estimations for the largest scenario (15 VMs) are as follows: the upper bound is overestimated by 16%, and the lower bound is underestimated by 34%. Our

Approach	6 VMs	9 VMs	12 VMs	15 VMs
On-premise Actual	1511	1053	981	796
ARIA Upper Bound	1765	1287	1059	924
ARIA Lower Bound	1379	892	660	523
100 Mbps weak link (3 VMs on-premise)				
Hybrid Actual	4886	3764	3794	3330
Hybrid Upper Bound	5490	4006	4764	4064
Hybrid Lower Bound	4675	3607	3140	3102
1 Gbps weak link (3 VMs on-premise)				
Hybrid Actual	1585	1130	1059	892
Hybrid Upper Bound	1831	1256	1351	1006
Hybrid Lower Bound	1509	1037	850	767

(a) Total actual runtime (s) vs. predicted runtime (s) expressed as upper bound and lower bound

TABLE 4: PageRank: Balanced example of an iterative MapReduce application with complex (multi-job) iterations

approach, on the other hand, has the following accuracy with respect to hybrid actual: for 15 VMs and 100 Mbps weak link, the upper bound is overestimated by 22%, and the lower bound is underestimated by 7%. For the 1 Gbps case, the upper bound is overestimated by 12%, and the lower bound is underestimated by 16%.

The average estimations are depicted in Figure 6. As can be observed, when the application is balanced and exhibits both map-intensive and reduce-intensive behavior, the weak link is under I/O pressure, especially in the 100 Mbps case. Therefore, the average ARIA estimation has an error of almost 80% with respect to hybrid actual. In the 1 Gbps case, the error is smaller, at 18%. Nevertheless, our approach exhibits again an error that is an order of magnitude smaller: less than 8% for the 100 Mbps case and less than 1% for the 1 Gbps case.

Table 4b shows finer-grained details about the accuracy of the map and reduce phases. For PageRank, the main source of the overall error seems to be the reduce phase estimation, which clearly overshadows the map phase error. This holds for both 100 Mbps and 1 Gbps. Our approach has a good accuracy for both phases regardless of the weak link capacity.

The final application we study is Connected Components, which emphasizes the reduce phase. For the scale of our experiments, we generated 300 MB of input data (representing interactions in a social network). The application runs for 9 complex iterations, each of which is composed by 3 MapReduce jobs and includes additionally one final job. Connected Components runs in the baseline case (3 on-premise VMs) for 1621 s. Doubling the number of on-premise VMs leads to an actual runtime that is 42% smaller, thereby showing good scalability potential.

Table 5a shows the upper and lower bounds of the estimations. As can be observed, there is a large error for ARIA versus on-premise actual, especially for the lower bound (more than 300%). In this case, our approach has almost an order of magnitude better approximation (35%) than does hybrid actual.

The average estimations depicted in Figure 7 exhibit a similar trend: in the 100 Mbps case, ARIA has an error of 64% compared with hybrid actual where our approach reduces this to 11%. In the 1 Gbps case, ARIA has a 21% error that our approach reduces to less than 2%.

Prediction	Accuracy vs. Actual	6 VMs	9 VMs	12 VMs	15 VMs
100 Mbps weak link (3 VMs on-premise)					
ARIA Avg.	(%) Map Error	-1.2	-6.6	-16.1	-13.1
	(%) Red. Error	-77.3	-79.8	-84.5	-85.3
	(%) Total Error	-67.8	-71.0	-77.4	-78.3
Hybrid Avg.	(%) Map Error	-0.1	-6.7	-7.4	-14.2
	(%) Red. Error	4.6	2.2	5.6	10.0
	(%) Total Error	4.0	1.1	4.2	7.6
1 Gbps weak link (3 VMs on-premise)					
ARIA Avg.	(%) Map Error	8.6	2.5	-11.6	-8.6
	(%) Red. Error	-5.8	-6.9	-22.8	-24.2
	(%) Total Error	-0.8	-3.6	-18.9	-18.9
Hybrid Avg.	(%) Map Error	5.2	-4.0	-9.0	-14.0
	(%) Red. Error	5.4	4.7	11.2	6.7
	(%) Total Error	5.4	1.5	3.9	-0.7

(b) Accuracy of the average prediction (between lower and upper bound) vs. the hybrid actual runtime broken down by phase

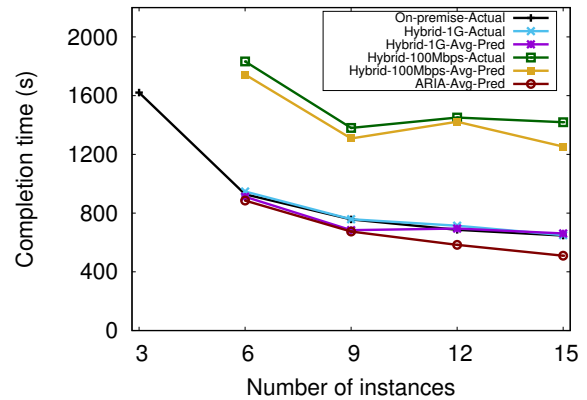


Fig. 7: Connected Components runtime estimation

Table 5b shows finer-grained details about the accuracy of the map and reduce phases. Since the reduce phase dominates, large errors in the reduce phase translate to low overall accuracy. ARIA exhibits these large errors in the reduce phase whereas our approach does not, thereby explaining the better overall accuracy.

6 CONCLUSIONS

This paper addresses the problem of how to estimate the runtime of iterative MapReduce applications in hybrid cloud-bursting scenarios where on-premise and off-premise VMs that host a MapReduce environment need to communicate over a weak link. Such runtime estimations are a critical tool in deciding whether the pay-as-you-go cost of cloud bursting justifies the expected speed-up.

To address this problem, we proposed a methodology that combines analytical modeling with synthetic benchmarking to estimate the time to solution specifically for a hybrid setup, where the weak link has a decisive impact on both the map and reduce phases. We illustrated our proposal for the MapReduce runtime; however, the principles are generic and can be applied to other big data analytics runtimes (e.g., Spark [28]) that support iterative computations.

We have demonstrated the benefits for our proposal from multiple angles using a mix of map-intensive, reduce-

Approach	6 VMs	9 VMs	12 VMs	15 VMs
On-premise Actual	928	757	686	647
ARIA Upper Bound	1178	972	883	810
ARIA Lower Bound	592	377	284	209
100 Mbps weak link (3 VMs on-premise)				
Hybrid Actual	1833	1380	1451	1419
Hybrid Upper Bound	2131	1585	1878	1450
Hybrid Lower Bound	1356	1031	967	1056
1 Gbps weak link (3 VMs on-premise)				
Hybrid Actual	947	758	714	650
Hybrid Upper Bound	1157	848	896	787
Hybrid Lower Bound	665	520	494	536

(a) Total actual runtime (s) vs. predicted runtime (s) expressed as upper bound and lower bound

Prediction	Precision vs. Actual	6 VMs	9 VMs	12 VMs	15 VMs
100 Mbps weak link (3 VMs on-premise)					
ARIA Avg.	(%) Map Error	-12.9	-15.9	-29.0	-35.1
	(%) Red. Error	-66.8	-66.5	-72.9	-76.8
	(%) Total Error	-51.7	-51.1	-59.8	-64.1
Hybrid Avg.	(%) Map Error	-6.0	-14.2	-11.8	-6.8
	(%) Red. Error	-4.5	-1.3	2.2	-13.8
	(%) Total Error	-4.9	-5.2	-2.0	-11.7
1 Gbps weak link (3 VMs on-premise)					
ARIA Avg.	(%) Map Error	-4.5	-8.0	-15.4	-21.6
	(%) Red. Error	-8.3	-13.6	-14.4	-21.9
	(%) Total Error	-6.5	-11.0	-18.3	-21.6
Hybrid Avg.	(%) Map Error	-2.8	-13.7	-3.8	5.7
	(%) Red. Error	-4.5	-5.2	7.0	-3.5
	(%) Total Error	-3.8	-9.8	-2.7	1.7

(b) Accuracy of the average prediction (between lower and upper bound) vs. the hybrid actual runtime broken down by phase

TABLE 5: Connected Components: Reduce-intensive example of an iterative MapReduce application

intensive, and balanced real-world iterative applications from standardized big data benchmarks that cover a broad spectrum of use cases. Specifically, we have shown that (1) the upper and lower estimation bounds of our approach against the hybrid baseline are significantly more accurate than the single-site counterparts against the on-premise baseline; (2) the average estimation of our approach is always within 1–10% error regardless of scale and up to one order of magnitude more accurate than single-site state-of-art against the hybrid baseline; (3) our approach shows consistent behavior and accurately estimates both the map and the reduce phases, indicating that the overall estimation was not obtained by accident through the accumulation of large errors during the map and reduce phases that cancel each other out. The trade-off for using our approach is the need for a one-time calibration phase; however, the results can be cached and reused for all subsequent applications, even by different users. Such an overhead is small considering that accurate estimations can bring such large benefits.

Encouraged by these results, we plan to explore an elastic approach where the number of off-premise VMs is not fixed but fluctuates to accommodate changing goals (e.g., results needed faster). Elastic performance is not yet well understood but is of critical practical importance for hybrid clouds. Furthermore, we plan to extend our proposal beyond MapReduce to predict performance for iterative computations using other big data analytics frameworks.

ACKNOWLEDGMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract DE-AC02-06CH11357, and the Spanish CICYT projects TIN2014-53495-R and TIN2017-82972-R.

REFERENCES

- [1] Cisco, "White paper: Cisco vni forecast and methodology, 2016."
- [2] T. Guo, U. Sharma, T. Wood, S. Sahu, and P. Shenoy, "Seagull: Intelligent cloud bursting for enterprise applications," in *USENIX ATC '12: Conference on Annual Technical Conference*, Berkeley, CA, USA, 2012, pp. 33–33.
- [3] F. J. Clemente-Castelló, B. Nicolae, R. Mayo, J. C. Fernández, and M. M. Rafique, "On exploiting data locality for iterative mapreduce applications in hybrid clouds," in *BDCAT '16: 3rd IEEE/ACM International Conference on Big Data Computing, Applications and Technologies*, Shanghai, China, 2016, pp. 118–122.
- [4] F. J. Clemente-Castelló, B. Nicolae, K. Katrinis, M. M. Rafique, R. Mayo, J. C. Fernández, and D. Loreti, "Enabling big data analytics in the hybrid cloud using iterative MapReduce," in *UCC '15: 8th IEEE/ACM International Conference on Utility and Cloud Computing*, Limassol, Cyprus, 2015, pp. 290–299.
- [5] T. White, *Hadoop: The Definitive Guide*. USA: Yahoo! Press, 2010.
- [6] F. J. Clemente-Castello, B. Nicolae, M. M. Rafique, R. Mayo, and J. C. Fernandez, "Evaluation of data locality strategies for hybrid cloud bursting of iterative mapreduce," in *CCGrid'17: 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, Madrid, Spain, 2017, pp. 181–185.
- [7] T. Gunarathne, T.-L. Wu, J. Qiu, and G. Fox, "MapReduce in the clouds for science," in *CloudCom '10: 20th IEEE Conference on Cloud Computing Technology and Science*, 2010, pp. 565–572.
- [8] X. Zhang, L. T. Yang, C. Liu, and J. Chen, "A scalable two-phase top-down specialization approach for data anonymization using MapReduce on cloud," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 2, pp. 363–373, 2014.
- [9] B. Nicolae, P. Riteau, and K. Keahey, "Bursting the cloud data bubble: Towards transparent storage elasticity in IaaS clouds," in *IPDPS '14: 28th IEEE International Parallel and Distributed Processing Symposium*, Phoenix, USA, 2014, pp. 135–144.
- [10] —, "Transparent Throughput Elasticity for IaaS Cloud Storage Using Guest-Side Block-Level Caching," in *UCC'14: 7th IEEE/ACM International Conference on Utility and Cloud Computing*, London, UK, 2014.
- [11] B. Nicolae, C. Costa, C. Misale, K. Katrinis, and Y. Park, "Leveraging adaptive I/O to optimize collective data shuffling patterns for big data analytics," *IEEE Transactions on Parallel and Distributed Systems*, 2017, to appear.
- [12] Y. Bu, B. Howe, M. Balazinska, and M. D. Ernst, "HaLoop: Efficient iterative data processing on large clusters," *Proc. VLDB Endow.*, vol. 3, no. 1-2, pp. 285–296, 2010.
- [13] Y. Zhang, Q. Gao, L. Gao, and C. Wang, "iMapReduce: A distributed computing framework for iterative computation," *Journal of Grid Computing*, vol. 10, no. 1, pp. 47–68, 2012.
- [14] B. Nicolae, P. Riteau, and K. Keahey, "Towards transparent throughput elasticity for IaaS cloud storage: Exploring the benefits of adaptive block-level caching," *International Journal of Distributed Systems and Technologies*, vol. 6, no. 4, pp. 21–44, 2015.
- [15] F. Tian and K. Chen, "Towards optimal resource provisioning for running MapReduce programs in public clouds," in *CLOUD '11: IEEE International Conference on Cloud Computing*, Washington, DC, USA, 2011, pp. 155–162.
- [16] K. Chen, J. Powers, S. Guo, and F. Tian, "CRESP: Towards optimal resource provisioning for MapReduce computing in public clouds," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 6, pp. 1403–1412, 2014.
- [17] P. Lama and X. Zhou, "AROMA: Automated resource allocation and configuration of MapReduce environment in the cloud," in *ICAC '12: 9th International Conference on Autonomic Computing*, New York, NY, USA, 2012, pp. 63–72.
- [18] H. Herodotou, H. Lim, G. Luo, N. Borisov, L. Dong, F. B. Cetin, and S. Babu, "Starfish: A self-tuning system for big data analytics," in *CRID '11: 5th Biennial Conference on Innovative Data Systems Research*, California, USA, 2011, pp. 261–272.

- [19] A. Verma, L. Cherkasova, and R. H. Campbell, "ARIA: Automatic Resource Inference and Allocation for MapReduce Environments," in *ICAC '11: 8th ACM International Conference on Autonomic Computing*, Karlsruhe, Germany, 2011.
- [20] A. Verma and R. H. Cherkasova, Ludmilaand Campbell, "Resource provisioning framework for MapReduce jobs with performance goals," in *Middleware '11: 12th ACM/IFIP/USENIX International Middleware Conference*, Lisbon, Portugal, 2011, pp. 165–186.
- [21] H. Herodotou, "Hadoop Performance Models," CS-2011-05, Computer Science Department Duke University, Tech. Rep., 2011.
- [22] Z. Zhang, L. Cherkasova, and B. T. Loo, "Benchmarking approach for designing a mapreduce performance model," in *ICPE '13: 4th ACM/SPEC International Conference on Performance Engineering*, 2013, pp. 253–258.
- [23] —, "Performance modeling of mapreduce jobs in heterogeneous cloud environments," in *CLOUD '13: 6th IEEE International Conference on Cloud Computing*, Washington, DC, USA, 2013, pp. 839–846.
- [24] M. Zaharia, A. Konwinski, A. D. Joseph, R. Katz, and I. Stoica, "Improving mapreduce performance in heterogeneous environments," in *OSDI '08: 8th USENIX Conference on Operating Systems Design and Implementation*, 2008, pp. 29–42.
- [25] F. Ahmad, S. T. Chakradhar, A. Raghunathan, and T. N. Vijaykumar, "Tarazu: Optimizing mapreduce on heterogeneous clusters," in *ASPLOS '12: 17th International Conference on Architectural Support for Programming Languages and Operating Systems*, 2012, pp. 61–74.
- [26] J. Polo, D. Carrera, Y. Becerra, V. Beltran, J. Torres, and E. Ayguadé, "Performance management of accelerated MapReduce workloads in heterogeneous clusters," 2010.
- [27] K. Shvachko, H. Huang, S. Radia, and R. Chansler, "The Hadoop distributed file system," in *MSST '10: 26th IEEE Symposium on Massive Storage Systems and Technologies*, 2010.
- [28] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica, "Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing," in *NSDI'12: 9th USENIX Conference on Networked Systems Design and Implementation*, San Jose, USA, 2012, pp. 2:1–2:14.
- [29] G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski, "Pregel: A system for large-scale graph processing," in *SIGMOD'10: The 2010 ACM SIGMOD International Conference on Management of Data*, Indianapolis, IN, USA, 2010, pp. 135–146.
- [30] Apache, "Apache hadoop rumen," <https://hadoop.apache.org/docs/r1.2.1/rumen.html>, accessed: 13-02-2018.
- [31] S. Godard, "Sysstat utilities for the Linux OS," <http://sebastien.godard.pagesperso-orange.fr/>, accessed: 13-02-2018.
- [32] A. Alexandrov, R. Bergmann, S. Ewen, J. Freytag, F. Hueske, A. Heise, O. Kao, M. Leich, U. Leser, V. Markl, F. Naumann, M. Peters, A. Rheinländer, M. J. Sax, S. Schelter, M. Höger, K. Tzoumas, and D. Warneke, "The stratosphere platform for big data analytics," *Vldb J.*, vol. 23, no. 6, pp. 939–964, 2014.
- [33] S. Ahn and S. Park, "An analytical approach to evaluation of ssd effects under mapreduce workloads," in *Journal of Semiconductor Technology and Science*, vol. 15, 10 2015, pp. 511–518.
- [34] S. H. Mohamed, T. E. H. El-Gorashi, and J. M. H. Elmighani, "On the energy efficiency of mapreduce shuffling operations in data centers," in *ICTON'17: 19th International Conference on Transparent Optical Networks*, Girona, Spain, 2017, pp. 1–5.
- [35] H.-H. Bock, "Clustering methods: A history of K-Means algorithms," in *Selected Contributions in Data Analysis and Classification*, 2007, pp. 161–172.
- [36] W. Zhao, H. Ma, and Q. He, "Parallel K-Means clustering based on MapReduce," in *CloudCom '09: 1st International Conference on Cloud Computing*, Beijing, China, 2009.
- [37] S. Huang, J. Huang, J. Dai, T. Xie, and B. Huang, "The HiBench benchmark suite: Characterization of the MapReduce-based data analysis," in *ICDEW '10: 26th IEEE International Conference on Data Engineering Workshops*, 2010, pp. 41–51.
- [38] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Comput. Netw. ISDN Syst.*, vol. 30, no. 1-7, pp. 107–117, 1998.
- [39] T. Seidl, B. Boden, and S. Fries, *CC-MR – Finding Connected Components in Huge Graphs with MapReduce*. Berlin, Heidelberg: Springer, 2012, pp. 458–473.
- [40] L. Wang, J. Zhan, C. Luo, Y. Zhu, Q. Yang, Y. He, W. Gao, Z. Jia, Y. Shi, S. Zhang *et al.*, "BigDataBench: A big data benchmark suite from internet services," in *HPCA '14: 20th IEEE International*

Symposium on High Performance Computer Architecture, 2014, pp. 488–499.




Francisco J. Clemente-Castelló received his B.Sc. in computer science in 2011, M.Sc. in intelligent systems in 2012, and Ph.D. in 2017 from the Department of Computer Science and Engineering, Jaume I University, Castellón, Spain. His main research interests include high-performance computing platforms, virtualization, cloud computing technologies, and big data analysis.



Bogdan Nicolae is a computer scientist at Argonne National Laboratory. He specializes in scalable storage, data management, and fault tolerance for large-scale distributed systems, with a focus on cloud computing and high-performance architectures. He holds a Ph.D. from the University of Rennes 1, France and a Dipl. Eng. degree from Politehnica University Bucharest, Romania. He has (co)authored numerous papers in the areas of scalable I/O, storage elasticity and virtualization, data and metadata decentralization and availability, multiversioning, checkpoint-restart, and live migration.



Rafael Mayo is an associate professor in the Computer Science and Engineering Department in Jaume I University, Spain. He received his B.Sc. and Ph.D. from Polytechnic Valencia University, Spain. His research interests include the optimization of numerical algorithms for general processors as well as for specific hardware (GPUs) and their parallelization on both message-passing parallel systems and shared-memory multiprocessors (SMPs, CCNUMA multiprocessors, and multicore processors). More recently, he has been involved in several research efforts on high-performance computing energy-aware systems and cloud computing.



Juan Carlos Fernández is an associate professor at Jaume I University of Castellón, Spain. He received his B.Sc. and Ph.D. from Polytechnic Valencia University, Spain. His research interests include the following topics: control algorithms in robot manipulators using parallel computing, parallel implementations of video encoder, energy saving on high-performance computing platforms, and cloud computing.