

A data-aware scheduling strategy for executing large-scale distributed workflows

SALVATORE GIAMPÀ¹, LORIS BELCASTRO¹, FABRIZIO MAROZZO^{1,2}, DOMENICO TALIA^{1,2}
and PAOLO TRUNFIO^{1,2}

¹DIMES - University of Calabria, Rende, Italy
²DtoK Lab Srl, Rende, Italy

Corresponding author: F. Marozzo (e-mail: fmarozzo@dimes.unical.it).

ABSTRACT Task scheduling is a crucial key component for the efficient execution of data-intensive applications on distributed environments, by which many machines must be coordinated to reduce execution times and bandwidth consumption. This paper presents ADAGE, a data-aware scheduler designed to efficiently execute data-intensive workflows in large-scale computers. The proposed scheduler is based on three key features: *i) critical path analysis*, for discovering the critical tasks of a workflow and reducing data transferring between nodes; *ii) work giving*, a new dynamic planning strategy for migrating tasks from overloaded to unloaded nodes; and *iii) task replication*, which executes task replicas on different nodes for improving both execution time and fault tolerance. Experiments performed on a distributed computing environment composed of up to 1,024 processing nodes show that ADAGE achieves better performances than existing scheduling systems, obtaining an average reduction of up to 66% in execution time.

INDEX TERMS Data-aware scheduler, workflow scheduling, distributed workflows, parallel programming, distributed computing, exascale computing

I. INTRODUCTION

The term Exascale refers to the capabilities of future computing systems, still to be implemented, which should be capable of calculating at least one exaFLOPS (i.e., 10^{18} FLOPS), far exceeding the most advanced existing computing systems (about 10^{15} FLOPS). To reach the Exascale size, other than new hardware solutions, it is required to define new programming models, languages and algorithms that combine abstraction with both scalability and performance [1]. Hybrid models (based on shared/distributed memory) and communication mechanisms based on data locality and grouping are currently investigated as promising approaches. Parallel applications running on Exascale systems will require to control a high number of tasks running on a very large set of computing resources [2]. Such applications will need to avoid or limit synchronization, use less communication and remote memory, and handle software and hardware faults that can occur. In order to achieve such computational speeds, more and more novel solutions are being proposed with the aim of harnessing the computational power of a large set of

machines operating in parallel [3].

The problem of coordinating many machines in a complex distributed system is widely represented as a task scheduling problem. Task scheduling has long been recognized as a NP-Hard problem, which represents a major challenge for researchers, especially if the scheduling is performed dynamically and in real time (as required by modern systems). The scheduling aims at identifying the most suitable resources for executing the workloads on time and optimizing resource utilization. In particular, it must also allow for many tasks to run simultaneously and for the exchange of large amounts of data. This paper presents ADAGE (A Data-aware scheduler based on critical path, work assignment, and replication), a data-aware scheduler designed to efficiently execute data-intensive workflows in a large-scale computer network. The proposed scheduler is based on three key features: *i) critical path analysis*, for discovering the critical tasks of a workflow and reducing data transferring between nodes; *ii) work giving*, a new dynamic planning strategy for migrating tasks from overloaded to unloaded nodes; and *iii) task replication*,

which executes task replicas on different nodes for improving both execution times and fault tolerance.

ADAGE is composed of the following components: *i*) *Decision Maker (DM)*, which runs on each node and assigns the tasks to the current node or remote nodes; *ii*) *Local Ready Queue (LRQ)*, which contains the tasks that are ready to be executed, sorted by execution priority; *iii*) *Load Balancer (LB)*, which moves tasks from the LRQ to the least loaded neighboring nodes whenever the current node is overloaded; and *iv*) *Validator*, which checks and updates the status of the tasks in LRQ. These components work together to effectively execute the submitted workflow, which is composed of many tasks with dependencies, on a pool of computing nodes in a distributed platform.

To evaluate our strategy, we carried out several experiments on different workflows by varying both the number of tasks and computing nodes. In our evaluations, we compared the designed scheduling strategy with two state-of-the-art systems, i.e., Matrix [4] and Albatross [5]. In particular, five existing workflows (i.e., *CyberShake*, *Epigenomics*, *Inspirational*, *Montage*, *Sipht*), which can generate up to 10,000 tasks, were evaluated. Experiments performed on a HPC distributed system composed of 1,024 computing nodes show that ADAGE achieves better performance than existing scheduling systems, obtaining an average reduction of up to 66% in execution time. For the purpose of using the code of our scheduler and allowing the reproducibility of the experiments, an open-source version of ADAGE is available at <https://github.com/SCALabUnical/ADAGE>.

Compared to existing techniques, our scheduler includes the following innovative aspects: *i*) it combines both static and dynamic planning strategies for reducing execution time of data-intensive workflows; *ii*) it exploits a novel algorithm for moving tasks from overloaded to unloaded nodes at runtime; *iii*) it takes advantage of task replication on different nodes to improve both execution times and fault tolerance.

The remainder of the paper is organized as follows. Section II discusses related work. Section III describes the proposed scheduling strategy. Section IV shows the experimental results and Section V concludes the paper.

II. RELATED WORK

With the increasing popularity of data-intensive workflows, several research projects have been carried out to define data-aware scheduling algorithms [6] [7] [8] aiming at improving scalability, energy efficiency and execution performance. In particular, due the imminent implementation of Exascale systems, task scheduling for massively parallel applications has become an important and strategic research area [3]. In particular, several algorithms and systems have been proposed to cope with the needs of large scale data-intensive applications, exploiting both static and dynamic scheduling [9] [10] [11].

Kosar et al. [12] proposed a data scheduler, namely *Stork*, that allows for planning data allocation and data transfer among computing nodes in a network. In particular, *Stork* uses the *ClassAd* [17] language to represent data manage-

ment tasks, while computational workflows are executed using both *Pegasus* [18] for data-aware planning and *HTCondor DAGMan* [19] for managing task dependencies. *Stork* exposes four main data scheduling strategies: first fit, largest fit, smallest fit and best fit. The first three strategies are heuristics, while the fourth one is an exact algorithm that discovers the best data allocation using a greedy approach.

Wei et al. [13] proposed a data-aware scheduling strategy obtained as the combination of two existing systems: LSF (Load Sharing Facility) [20] and *GFarm* [21]. LSF is a job scheduler expressively designed for HPC systems that exposes a set of scheduling tools for managing global workloads and resources. *GFarm* is a distributed file system that is designed for data sharing in large clusters. The proposed strategy exploits data location information retrieved from *GFarm* for evaluating data affinity of tasks and automatically transfer data among nodes.

Acevedo et al. [14] proposed a data-aware scheduling algorithm based on a variant of the critical path algorithm [22], named Critical Path File Location (CPFL). The algorithm is designed to schedule workflow tasks by declaring inter-task and data dependencies. It also allows to execute an application composed of multiple workflows by merging them in a single meta-workflow. The scheduler exploits a pre-scheduling stage to establish where data should be allocated. Then, the critical path algorithm is used to assign a priority value to each task of the meta-workflow. Subsequently, the scheduler manages each task using priority and assigns it to a computing node based on its data dependencies.

Marozzo et al. [15] proposed a composition of two systems, the Data Mining Cloud Framework (DMCF) [23] and *Hercules* [24], to obtain a data-aware workflow scheduling for Cloud environments. DMCF allows to process and schedule workflow tasks, while *Hercules* manages temporary files generated during computation. The scheduling strategy is inspired by the one proposed in [25], but it uses a new local queue on the executor node, called *locallyActivatedTask*, to obtain data-awareness. For each node, the scheduler selects the best task to run, choosing it from the global or local task queue. In particular, the scheduler tries to execute the task whose dependencies have been resolved and for which the current node is the best concerning data allocation.

MATRIX (MAny-Task computing execution fabRIc at eXascale) [16] is a system that implements a data-aware scheduling strategy based on work stealing [26]. It extends the classical work stealing strategy to support data-awareness by maintaining information about data dependencies of scheduled tasks. The system consists of three entities: client, scheduler and executor. The network nodes are fully connected, which means that they can communicate each other. On each node, three basic components run: executor, scheduler, and *ZHT (Zero-hop distributed HashTable)* server [27]. In particular, the last component allows to implement a shared *DKVS (Distributed Key-Value Store)* that stores information about tasks, including data dependencies and data locality. *MATRIX* exploits three local queues to

Reference	Metadata Management System	Storage system	Language	Scheduler features	Evaluated metrics
Kosar et al. [12]	HTCondor DAGMan	Pegasus	Go	Stork HTCondor-G	Transfer speed
Wei et al. [13]	LSF	GFarm	C\C++	LSF plugin	-
Acevedo et al. [14]	DHT	Distributed/Local filesystem	Java	CPFL	Execution time
DMCF [15]	Microsoft Azure	Hercules Microsoft Azure	Javascript VL4Cloud	In-memory temporary file	I/O operations, Execution time
MATRIX [16]	ZHT	ZHT	C\C++	Work Stealing	Throughput
Albatross [5]	ZHT	HDFS	C\C++	Late Binding	Latency, Throughput
<i>ADAGE</i>	<i>DHT</i>	<i>Distributed/Local file system</i>	<i>Java</i>	<i>Critical Path Work Giving Backup tasks</i>	<i>Execution time Throughput</i>

TABLE 1: Comparison with main related systems.

manage tasks, which contain respectively: tasks that are not ready to run; strictly data-dependent tasks that read much data from well-defined nodes; and not strictly data-dependent tasks that read some temporary data.

Albatross [5] is a system that improves some features of MATRIX. For example, it enhances fault tolerance by replacing the local queue containing not ready tasks with Fabriq [28]. Fabriq is a distributed message queue (DMQ) that runs on top of a distributed hash table, which prevents losing tasks when a node fails. Albatross assign tasks to nodes by using a *late-binding* technique. Specifically: *i*) when a task becomes ready (i.e., all its dependencies are solved), the load balancer pulls it from the DMQ and tries to send it to the best node according to data locality; *ii*) if the remote node is overloaded or data are local, the task is assigned to the local queue of the current node; *iii*) when the task is pulled from a local queue, the load balancer tries again to send it to the best node and, if the assignment is not possible, the task is executed on the current node.

Table 1 shows a comparison among the referred related works. For each work, the table reports the metadata management and storage systems, implementation language, features of the scheduler, and performance metrics that have been evaluated. Differently from existing techniques, ADAGE combines both static and dynamic planning strategies for improving the execution performances of data-intensive workflows. In particular, a static planning strategy, based on the critical path algorithm, is used to optimally assign tasks to the nodes during the workflow submission. Then, a novel dynamic strategy, named *work-giving*, is used by overloaded nodes for assigning tasks to other nodes. Furthermore, ADAGE exploits task replication on different nodes to improve both execution times and fault tolerance.

III. PROPOSED SCHEDULER

ADAGE is a new data-aware scheduler that exploits data locality to reduce data movement among nodes and improve the execution time of data-intensive workflows. To reach this

goal, ADAGE combines both static and dynamic planning strategies.

The static planning strategy is based on the *critical path* algorithm [22], which permits to find the *critical tasks* of a workflow, i.e. tasks that cannot be delayed without delaying the execution of the entire workflow. Starting from the knowledge of the critical path, our strategy minimizes data movement and memory latency by executing a task on the node that holds the largest amount of input data.

A dynamic planning strategy is used for assigning tasks to computing nodes at runtime. We designed a new dynamic planning strategy, called *work-giving*, which is used for migrating tasks from overloaded to unloaded nodes. Specifically, if a node is overloaded, it tries to send some of its tasks to unloaded nodes in its neighborhood. Such behavior differs from the work-stealing approach, in which an entity runs on the unloaded nodes and, during the computation, searches and steals tasks from the overloaded ones. It should be noted that the stealing process is activated many times and in many nodes. This behavior can limit the scalability in large-scale computing systems (such as Exascale computers), where the unloaded nodes are usually much more than the overloaded ones. Additionally, each unloaded node competes with the others to steal tasks, which can lead to a highly random distribution of tasks in the system. Differently, the *work-giving* strategy is executed on a much smaller number of nodes, which improves the system scalability. In addition, this approach limits the random distribution of the tasks by allowing an overloaded node to assign tasks to a limited number of nodes in its neighborhood.

For increasing the application reliability and finishing computation faster, ADAGE exploits *task replication* to execute speculative copies of tasks on different nodes. As stated in [29], the use of task replicas (also called backup tasks) is essential to significantly reduce the completion time of large workflow applications. In fact, some computing nodes may take an unusually long time to complete some tasks (e.g., due to overhead or hardware/software issues), negatively

affecting the completion time of the entire application. This mechanism marks a task as completed when the primary or a replica execution ends.

More details on the architecture, metadata and algorithms exploited by ADAGE are provided in the following sections.

A. ARCHITECTURE

The software structure of ADAGE consists of the following macro-components:

- *Client*: given a workflow composed of several tasks, it executes the critical path algorithm to pre-assign the tasks to the nodes and calculates the priority for each of them.
- *Distributed Hash Table (DHT)*: it stores all the necessary information about tasks, such as the running state, parent tasks, and actual number of replicas.
- *Distributed Message Queue (DMQ)*: it stores the identifiers of tasks waiting to be executed by a processing node.
- *Scheduler*: it executes a dynamic scheduling strategy, named *work-giving*, which is discussed in Section III.

An instance of the scheduler runs on each node of the system. Specifically, such a scheduler instance is composed of the following components:

- *Decision Maker (DM)*: it statically assigns tasks to the current node or remote nodes.
- *Local Ready Queue (LRQ)*: it contains the *ready tasks*, which are tasks whose dependencies are solved; such tasks are sorted by the execution priority calculated with the critical path algorithm.
- *Load Balancer (LB)*: when the current node is overloaded, the LB selects and sends some tasks from the LRQ to less loaded neighbor nodes.
- *Validator*: it checks the completion of tasks in the LRQ; if a task is completed, the Validator removes it from the queue; otherwise, if the task execution is still pending, the heartbeat is updated and stored in the DHT.

Figure 1 shows the block diagram and execution flow of the scheduling system. Each node can dispose of one or more Executors, which pull ready tasks from the LRQ. If the LRQ is empty, the execution of the Decision Maker is triggered.

As shown in Figure 1, the client starts storing task metadata in the DHT (1) and tasks in the DMQ (2) (also specifying the preferred execution node for each of them). In particular, the client executes the critical path algorithm to find an optimal planning solution for assigning the tasks to nodes. On each node, an instance of the Decision Maker (DM) takes one or more tasks from the DMQ (3) and decides on which node they have to be executed. Specifically, if a task has been assigned by the client to the current node, it is put in the Local Ready Queue (LRQ) (4); otherwise, such a task is sent to another node, which is chosen based on data locality (5), and inserted in the LRQ (6).

The Executors get the tasks with the highest priorities from the LRQ (7) for running them, while the Load Balancer (LB)

gets those with the lowest priorities (8). In the latter case, the tasks are replicated and sent to some neighbor nodes that are less charged than the current one (9). In particular, each replicated task is inserted in the LRQ of the chosen neighbor node (10). The maximum number of replicas for a task is specified by the client during the workflow submission and stored in the DHT.

Only the tasks that are assigned to the current node can be replicated by the LB. In fact, as shown in Figure 1, the LRQ is logically split in two parts so as to distinguish between the tasks that have been assigned to the current node and replicas that have been received from neighbors. More details about the different components of the scheduling system are provided in the following subsections.

B. TASK METADATA AND STATES

Table 2 reports the main metadata of a task. In particular, the field *state* represents the current state of the task, which can take one of the values reported in Table 3.

TABLE 2: The task metadata stored in the DHT.

<i>Metadata</i>	<i>Description</i>
ID	Task identifier
state	[waiting, ready, stage-in, running, failed, stage-out, complete]
heart-beat	timestamp of last heartbeat from the nodes having the task
children	set of children tasks
parents	number of parents the task is waiting for
in-list	names of input files the task reads
out-list	names of output files produced by the task
node-list	list of nodes having the task
num-backup	maximum number of backup nodes, given by the client
state-history	timestamps of task state changes (e.g., [{ts1:waiting},{ts2:ready},...])

TABLE 3: The possible states of a task.

<i>State</i>	<i>Description</i>
<i>waiting</i>	the task is waiting termination of parent tasks
<i>ready</i>	the task is ready to execute, because its parents are complete
<i>stage-in</i>	the task is ready to run, but it is waiting availability of input data
<i>running</i>	the task is running
<i>failed</i>	the task failed for some reason
<i>stage-out</i>	the task is complete, its output data are being written in the storage
<i>complete</i>	the task is complete, its output data are available from the storage

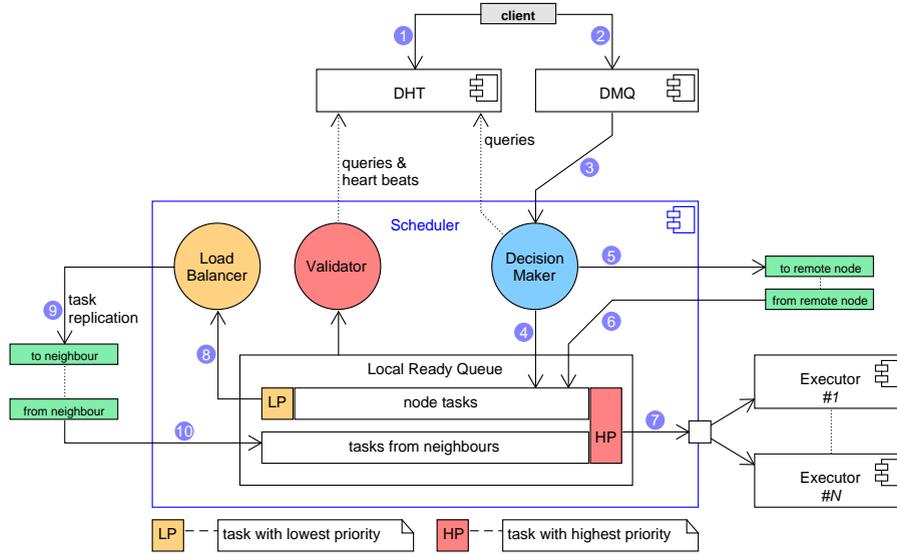


FIGURE 1: Scheduler block diagram and execution flow.

Figure 2 shows the state diagram describing the life cycle of a task. When it is submitted by the client, a task is in the *waiting* state, which means the task is waiting for termination of some parent task. The field *parents* contains the number of parents the task is waiting for. Such a number is decreased every time a parent task terminates.

The task switches to the *ready* state when all its parents have been completed successfully. In such a state, the node can start preparing the task for execution, getting the needed data. When this happens, the task goes in the *stage-in* state. Once the stage-in process is completed, the task is executed, passing in the *running* state. If the task fails its execution due to a self-generated error (e.g., a programming error or an unhandled exception), it goes in the *failed* state and all its children fail in cascade. Alternatively, if the task successfully completes its execution, the scheduler starts to write the output data in the storage and the task is switched in the *stage-out* state. Finally, the task goes in the *complete* state when the output data have been fully stored.

C. DECISION MAKER AND EXECUTOR

When an Executor terminates its current work, it selects another task from the LRQ. If the queue is empty, it activates the Decision Maker (DM), which starts to load new tasks from the DMQ to the LRQ. Then, the Executor tries again to get a task from the LRQ and, if found, executes it; otherwise it activates the DM again. However, to limit the network overload due to subsequent calls to the DMQ, the Executor awaits a short time before activating the DM again. Listing 1 shows the pseudo-code of the Executor component.

```

1  node: the current node the Executor is running on
2  nodeSet: the set of all nodes in the system
3  executorInterval: time to await between task pulling
   → attempts
4
5  function Executor(node, nodeSet, executorInterval):
6    repeat:
7      if LRQ is empty:
8        invoke DecisionMaker(nodeSet)
9      end if
10
11     if LRQ is not empty:
12       task := get the next non-finished task from LRQ
13       run the task to its completion
14       set the "state" field of task to "complete"
15       add a new timestamp for the "complete" state to
   → the "state-history" field of task
16     foreach child in "children" field of task:
17       decrease the "parent" field of child
18     end foreach
19   else
20     await executorInterval
21   end if
22 until not interrupted
23 end function
24

```

Listing 1: Pseudo-code of the Executor component. It invokes the Decision Maker when no tasks are present in the Local Ready Queue.

After executing a task, the Executor updates metadata of the task and its children. In particular:

- the *state* field is set to *complete*;
- a new timestamp for the *complete* state is added to the *state-history* field;
- for each *child* in the *children* field, the *parents* field is decreased.

The Decision Maker performs an initial distribution of tasks to nodes based on data locality. In particular, it performs the following steps:

- 1) It checks the DMQ looking for non-finished tasks that have been assigned to the current node and whose

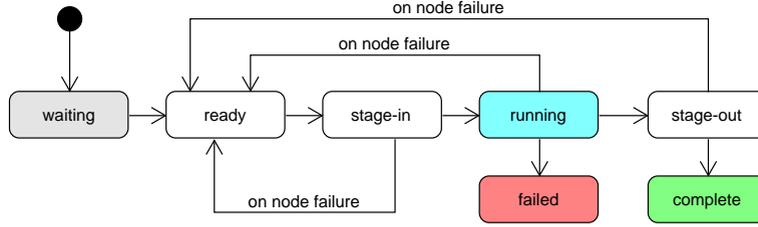


FIGURE 2: State diagram describing the life-cycle of a task.

heartbeats are not up-to-date. If some tasks are found, the DM inserts them in the LRQ and terminates its execution; otherwise, it proceeds to the second phase.

- 2) It scans the DMQ again looking for non-finished tasks that are not assigned to the current node and whose heartbeats are not up-to-date (i.e., the tasks that have been pulled from the DMQ but whose assignee node failed). If the DM gets any tasks matching such criteria, it decides where to send them. Specifically, if a task is assigned to the current node, it is inserted in the LRQ; otherwise, such a task is sent to another node in the system, which is optimally chosen based on data locality.

The DM aims at improving the fault tolerance of the system when the pre-assignment of tasks (made by the client) is no longer feasible (e.g., because some assignee node is failed or unreachable). Listing 2 shows the pseudo-code of the DM.

```

1 nodeSet: the set of all nodes in the system
2
3 function DecisionMaker(nodeSet):
4   taskFound := false
5   foreach task in DMQ do:
6     if task is not finished and task is not in waiting
7       ↪ state and task is preassigned to this node and
8       ↪ the last heartbeat of the task is too old then:
9         update the heartbeat of the task
10        insert task in the LRQ
11        taskFound := true
12    end if
13  end foreach
14  if taskFound is false then:
15    foreach task in DMQ do:
16      if task is not finished and task is not in waiting
17        ↪ state and task is not preassigned to this node
18        ↪ and the last the heartbeat of task is too old
19        ↪ then:
20        pre_node := the node to which the task has been
21        ↪ preassigned
22        if pre_node is not null and pre_node is online
23        ↪ then:
24          update heartbeat of the task
25          insert task in the LRQ of pre_node
26        else:
27          best_node := GetBestNode(task, nodeSet)
28          update heartbeat of the task
29          insert task in the LRQ of best_node
30        end if
31      end if
32    end foreach
33  end if
34 end function
  
```

Listing 2: Pseudo-code of the Decision Maker.

In particular, when a task cannot be executed by the assignee node, the DM sends it to another node for execution. The new assignee node is chosen using a heuristic that takes into account the location of the input data used by the task. Listing 3 shows the pseudo-code of the procedure used to find the best node for executing a task according to its input data locality. Similarly to what was proposed by Acevedo et al. [14], such a procedure aims at minimizing the total transfer time of all input data. In particular, the transfer time of a file is calculated as the ratio between file size and bandwidth of the node. Given a node, the total transfer time is calculated by considering all the input files, required by the task, that are owned by the node itself. Then, the node that grants the smallest total transfer time is chosen.

The time complexity of the Decision Maker function is $O(t * n)$, where t is the number of submitted tasks and n is the number of nodes in the computer network.

```

1 task: the task to search the best node for
2 nodeSet: the set of all nodes in the system
3
4 function GetBestNode(task, nodeSet)
5   best-node := null
6   best-data-transfer-time := 0
7   foreach node in nodeSet do:
8     if not task is already assigned to node then:
9       data-transfer-time := 0;
10      foreach file read by task do:
11        file-nodes := get all the nodes that contains a
12        ↪ replica of the file
13        if node is in file-nodes then:
14          data-transfer-time := data-transfer-time + (
15          ↪ size of file)/(bandwidth of node)
16        end if
17      end foreach
18      if best-node is null or data-transfer-time < best-
19      ↪ data-transfer-time then:
20        best-data-transfer-time := data-transfer-time
21        best-node := node
22      end if
23    end if
24  end foreach
25  return best-node
26 end function
  
```

Listing 3: Pseudo-code of the procedure used to find the best node for executing a task.

D. LOAD BALANCER

The Load Balancer (LB) is a periodic thread that monitors the workload of all nodes in the neighborhood in order to efficiently distribute tasks among them.

Listing 4 shows the pseudo-code of the LB component. In particular, after setting an initial waiting time $lbTime$, the LB gets the list of neighbor nodes, which can be retrieved by using a static or a dynamic approach. According to the static approach, the neighborhood does not change over time, while using the dynamic one it can be calculated many times. For example, MATRIX [4] uses a dynamic selection strategy that randomly chooses a number of neighbor nodes equal to the square root of the total number of nodes. Then, if the local node is the most overloaded one in the neighborhood, the LB sends half of the tasks to the less overloaded neighbor node. However, such operation can fail if: *i*) the local node is not the most overloaded one in the neighborhood; *ii*) a communication error happens; *iii*) the LRQ is empty. In such cases, $lbTime$ is doubled and the LB performs a new attempt after sleeping for that time. To avoid too long waits, $lbTime$ is doubled until it reaches a maximum allowed value. On the other hand, if the tasks are sent successfully, the waiting time is reset to the initial value (e.g., 1 millisecond). Before sending tasks to a neighbor, the LB replicates them and maintains the original copies in the LRQ. In such a way, the different replicas of a task compete each other to be executed first. The first replica that completes its execution determines the end of the task and, consequently, causes the termination of all other running replicas.

```

1 node: the local machine
2 maxTime: the maximum waiting time
3
4 function LoadBalancer(node, maxTime):
5   lbTime := 1
6   repeat:
7     success := false
8     neighbourhood := get neighborhood of this node
9     if node is the most overloaded one in neighbourhood
10      ↪ then:
11       neighbour := get the less overloaded node in the
12       ↪ neighbourhood
13       tasks := select half of the tasks assigned to this
14       ↪ node from the LRQ
15       if tasks is not empty then:
16         insert the replicas of all tasks in the LRQ of
17         ↪ the neighbor
18         lbTime := 1;
19         success := true
20       end if
21     end if
22     if not success and lbTime < maxTime:
23       double lbTime (lbTime := lbTime * 2)
24     end if
25     await lbTime
26   until not interrupted
27 end function

```

Listing 4: Pseudo-code of the Load Balancer.

As it can be observed from Listing 4, the time complexity of the Load Balancer function is $O(t + n)$, where t is the number of submitted tasks and n is the number of nodes in the network.

E. VALIDATOR

As explained in Section III-D, the Load Balancer can replicate the tasks on the neighborhood of the current node. In particular a replica of a task is inserted in the LRQ of a neighbor node. An additional component, namely the Val-

idator, monitors the LRQ looking for tasks that completed their execution. If any are found, the local replicas of such tasks are removed and not executed again. This operation can be accomplished by querying the DHT, which stores all the needed information about the tasks that are currently running in the system. The Validator also updates the heartbeats of both the tasks in the LRQ and those that are currently running. Listing 5 shows the pseudo-code of the Validator component.

```

1 runningTasks: the tasks currently running on this node
2 validationInterval: interval of time between two
3   ↪ activations of the Validator
4
5 function Validator(validationInterval):
6   repeat:
7     foreach task in runningTasks do:
8       update heartbeat of task
9     end foreach
10    foreach task in LRQ do:
11      if task is finished then:
12        remove task from LRQ
13      else:
14        update heartbeat of task
15      end if
16    end foreach
17    await validationInterval
18  until not interrupted
19 end function

```

Listing 5: Pseudo-code of the Validator.

On the single node, the Validator function has a time complexity that is linear in the number of tasks in the LRQ. Considering all the nodes, the total time complexity results to be $O(t)$, where t is the number of submitted tasks.

IV. CASE STUDIES AND EXPERIMENTS

We experimentally evaluated the performance of our scheduling strategy using WorkflowSim [30], a widely used toolkit for running distributed workflows, which allows to consider many aspects of the system, such as machine bandwidth, storage types (e.g., RAM, local disk or distributed file systems), hardware specifics (e.g., number of cores and clock speed) and power consumption. Each computing node in our tests has been configured with 4 CPU cores at 2,000 MIPS, 8 GB of RAM, 1 Gbps of bandwidth, and 1 TB of storage.

In our experiments, we used five existing workflows that are defined in [31]: *CyberShake*, *Epigenomics*, *Inspiral*, *Montage*, *Sipht*. To assess the effectiveness of our scheduling strategy, we compared it with two related systems: MATRIX [4] and Albatross [5]. In particular, we evaluated the execution time by varying both the number of nodes and tasks, the throughput as the number of completed tasks per second, and the distribution of the completed tasks over the execution time.

A. EXPERIMENT RESULTS

Figure 3 reports the elapsed execution time of the different scheduling systems when executing the five different workflows by varying the number of nodes from 1 to 1,024 (i.e., up to 4,096 cores). Specifically, each workflow has

been configured to spawn 1,000 tasks. Each test has been configured with the following parameters:

- number of replicas for each task: 2;
- heartbeat expiration period: 125 s;
- period between two subsequent activations of the Validator: 62.5s (i.e., half of the heartbeat expiration period).

As shown in Figure 3(a), for the Sipt application, our strategy results to be on average 21% and 13% faster than Albatross and MATRIX respectively. For other applications, the reduction of execution time ranges from 15% to 30% for Epigenomics (Figure 3(b)), from 20% to 31% for Inspiral (Figure 3(c)), from 18% to 66% for CyberShake (Figure 3(d)), and from 1% to 23% for Montage (Figure 3(e)).

Since our scheduler was designed to support Exascale applications, which can be composed of tens of thousands of tasks, we carried out additional experiments to evaluate the execution times when the number of tasks is increased up to 10,000. For the sake of brevity, we compared the performance of the three systems using only the Montage and CyberShake workflows. Figure 4 shows the execution time obtained by increasing the number of tasks up to 10,000 on 1,024 computing nodes. In particular, as the number of tasks increases, in both Figures 4(a) and 4(b) the execution time of ADAGE grows slower than that of the other two systems. This experiment demonstrates a greater ability of our scheduler to manage large computational resources.

Figure 5 illustrates the throughput of the different scheduling systems, which has been calculated as the number of completed tasks per second. As shown, ADAGE achieves significantly better results than the other systems. In particular, as the number of available nodes increases, the throughput of our scheduler considerably increases compared to that of the two other strategies, which demonstrates how ADAGE is particularly suitable for very large distributed computation systems. By executing the CyberShake workflow with 1,024 compute nodes, ADAGE obtains a throughput that is 83% and 1377% greater than that of Albatross and MATRIX respectively (Figure 5(a)). In the Montage workflow case, using 1,024 nodes, the throughput of ADAGE is 11% and 50% greater (Figure 5(b)) of Albatross and MATRIX respectively.

Figure 6 shows the distribution of the completed tasks over the execution time. The plotted results show that ADAGE achieves the peak of completed tasks much faster than the other two systems. This figure reports how the execution time of tasks is greatly reduced by using ADAGE, which means that computational resources are released in a shorter time with reference to the other two approaches.

Overall, the obtained results and the wide number of experiments carried out on different workflows demonstrated how the proposed scheduling strategy offers better performance than other existing systems. This is especially true when a very large number of nodes is used. This feature makes the proposed algorithm particularly interesting for supporting massive task execution in the upcoming Exascale systems.

V. CONCLUSION

In this paper we presented ADAGE, a new data-aware scheduling strategy for large distributed computation environments, such as the upcoming Exascale systems. Differently from existing techniques, ADAGE combines both static and dynamic planning strategies for improving the execution time of data-intensive workflows. In particular, it is based on three key features: *i) critical path analysis*, for discovering the critical tasks of a workflow and reducing data transferring between nodes; *ii) work giving*, a new dynamic planning strategy for migrating tasks from overloaded to unloaded nodes; and *iii) task replication*, which executes task replicas on different nodes for improving both execution times and fault tolerance. Experiments performed on a distributed environment composed of up to 1,024 compute nodes showed that ADAGE achieves better performances than existing techniques, obtaining a reduction of up to 66% in execution time. Moreover, as the number of available nodes increases, ADAGE outperformed the other techniques in terms of throughput, demonstrating that is particularly suitable for very large distributed computing systems.

In future work, additional research issues will be investigated. In particular, since our scheduler supports different workflows patterns (e.g., map-reduce, divide-and-conquer, pipeline), we plan to investigate its usability in combination with Apache Hadoop and Spark, which are widely used for developing and executing general-purpose high performance applications.

DATA AND CODE AVAILABILITY STATEMENT

For the purpose of using the code of our scheduler, an open-source version of ADAGE is available at <https://github.com/SCAalabUnical/ADAGE> along with some sample workflows and instructions for running experiments.

ACKNOWLEDGMENT

This work has been supported by the ASPIDE Project funded by the European Union's Horizon 2020 Research and Innovation Programme under grant agreement No 801091.

REFERENCES

- [1] Georges Da Costa, Thomas Fahringer, Juan-Antonio Rico-Gallego, Ivan Grasso, Atanas Hristov, Helen D. Karatza, Alexey Lastovetsky, Fabrizio Marozzo, Dana Petcu, Georgios L. Stavrinides, Domenico Talia, Paolo Trunfio, and Hrachya Astsatryan. Exascale machines require new programming paradigms and runtimes. *Supercomputing Frontiers and Innovations*, 2(2):6–27, 2015.
- [2] Loris Belcastro, Fabrizio Marozzo, and Domenico Talia. Programming models and systems for big data analysis. *International Journal of Parallel, Emergent and Distributed Systems*, 34:632–652, 2019.
- [3] Domenico Talia, Paolo Trunfio, Fabrizio Marozzo, Loris Belcastro, Javier Garcia Blas, David Del Rio, Philippe Couvee, Gael Goret, Lionel Vincent, Alberto Fernandez-Pena, Daniel Martin de Blas, Mirko Nardi, Teresa

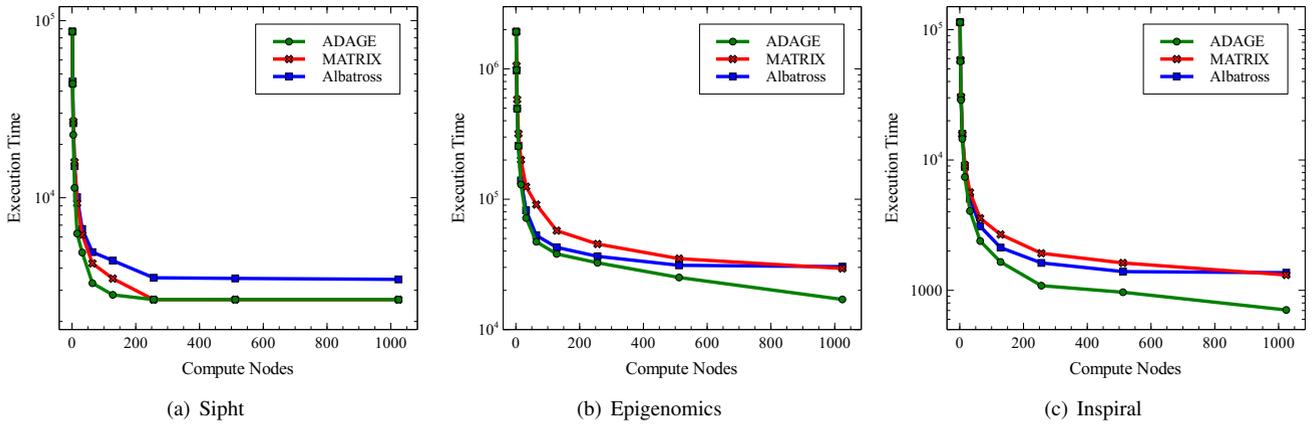


FIGURE 3: Comparison of the execution times by varying the number of compute nodes (five workflows and 1,000 tasks).

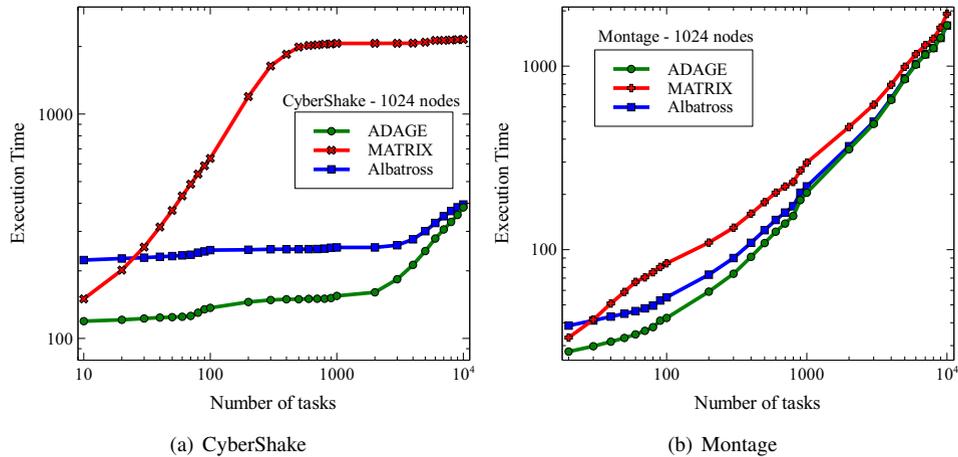


FIGURE 4: Comparison of the execution times by increasing the number of tasks up to 10,000 (two workflows and 1,024 compute nodes).

Pizzuti, Adrian Spataru, and Marek Justyna. A novel data-centric programming model for large-scale parallel systems. In Euro-Par 2019: Parallel Processing Workshops, Lecture Notes in Computer Science, pages 452–463, Gottingen, Germany, 26-30 August 2020. ISBN: 978-3-030-48339-5.

[4] Anupam Rajendran, Ke Wang, and Ioan Raicu. Matrix: Many-task computing execution fabric at exascale. 01 2013.

[5] Iman Sadooghi, Geet Kumar, Ke Wang, Dongfang Zhao, Tonglin Li, and Ioan Raicu. Albatross: An efficient cloud-enabled task scheduling and execution framework using distributed message queues. pages 11–20, 10 2016.

[6] Tevfik Kosar and Mehmet Balman. A new paradigm: Data-aware scheduling in grid computing. *Future Generation Computer Systems*, 25(4):406–413, 2009.

[7] Shivaram Venkataraman, Aurojit Panda, Ganesh Ananthanarayanan,

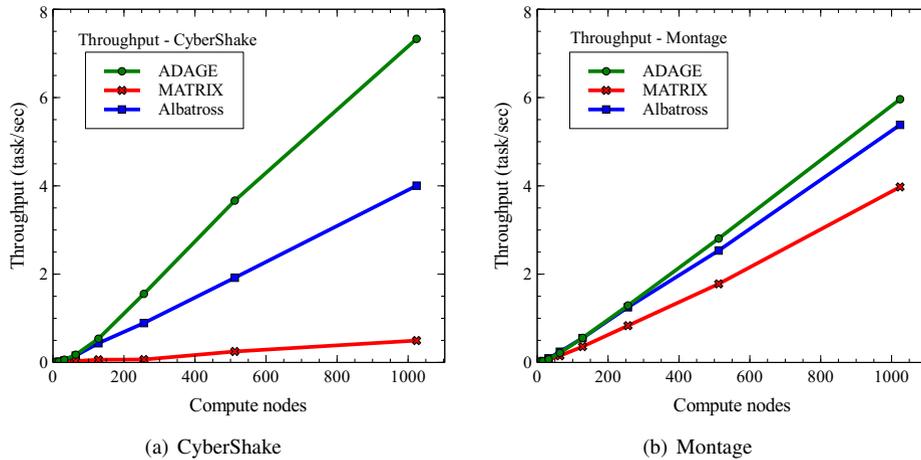


FIGURE 5: Comparison of number of tasks completed per second (throughput) vs the number of compute nodes.

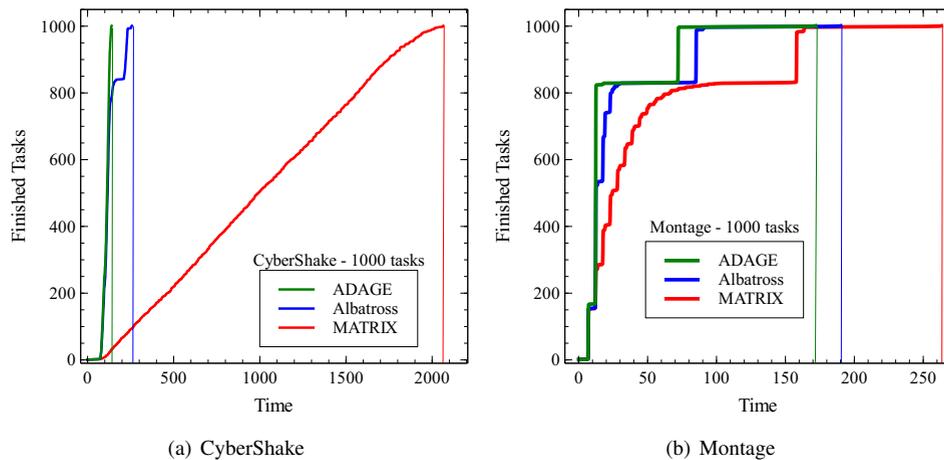


FIGURE 6: Comparison of task completion times using 1,024 compute nodes.

- Michael J Franklin, and Ion Stoica. The power of choice in data-aware cluster scheduling. In 11th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 14), pages 301–316, 2014.
- [8] Jiahui Jin, Junzhou Luo, Aibo Song, Fang Dong, and Runqun Xiong. Bar: An efficient data locality driven task scheduling algorithm for cloud computing. In 2011 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, pages 295–304. IEEE, 2011.
- [9] CL Philip Chen and Chun-Yang Zhang. Data-intensive applications, challenges, techniques and technologies: A survey on big data. *Information sciences*, 275:314–347, 2014.
- [10] Ji Liu, Esther Pacitti, Patrick Valduriez, and Marta Mattoso. A survey of data-intensive scientific workflow management. *Journal of Grid Computing*, 13(4):457–493, 2015.
- [11] Ke Wang, Kan Qiao, Iman Sadooghi, Xiaobing Zhou, Tonglin Li, Michael Lang, and Ioan Raicu. Load-balanced and locality-aware scheduling for data-intensive workloads at extreme scales. *Concurrency and Computation: Practice and Experience*, 28(1):70–94, 2016.
- [12] Tevfik Kosar, Mehmet Balman, Esmâ Yildirim, Sivakumar Kulasekaran, and Brandon Ross. Stork data scheduler: Mitigating the data bottleneck in e-science. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, 369:3254–67, 08 2011.
- [13] Xiaohui Wei, Wilfred Li, Osamu Tatebe, Gaochao Xu, Liang Hu, and Jiubin Ju. Implementing data aware scheduling in gfarm(r) using lsf(tm) scheduler plugin mechanism. pages 3–10, 01 2005.
- [14] César Acevedo, Porfidio Hernández, Antonio Espinosa, and Víctor Méndez. A critical path file location (cpfl) algorithm for data-aware multiwork-
flow scheduling on hpc clusters. *Future Generation Computer Systems*, 74, 04 2017.
- [15] Fabrizio Marozzo, Francisco Duro, Javier Garcia Blas, Jesus Carretero, Domenico Talia, and Paolo Trunfio. A data-aware scheduling strategy for workflow execution in clouds. *Concurrency and Computation: Practice and Experience*, page e4229, 08 2017.
- [16] Ke Wang, Xraobing Zhou, Tonglin Li, Dongfang Zhao, Michael Lang, and Ioan Raicu. Optimizing load balancing and data-locality with data-aware scheduling. In 2014 IEEE International Conference on Big Data (Big Data), pages 119–128. IEEE, 2014.
- [17] Rajesh Raman, Marvin Solomon, Miron Livny, and Alain Roy. The classads language. pages 255–270, 01 2004.
- [18] Ewa Deelman, Gurmeet Singh, Mei-Hui Su, James Blythe, Yolanda Gil, Carl Kesselman, Gaurang Mehta, Karan Vahi, G. Berriman, John Good, Anastasia Laity, and Daniel S. Katz. Pegasus: A framework for mapping complex scientific workflows onto distributed systems. *Scientific Programming*, 13:219–237, 01 2005.
- [19] Douglas Thain, Todd Tannenbaum, and Miron Livny. Distributed computing in practice: the condor experience. *Concurrency - Practice and Experience*, 17(2-4):323–356, 2005.
- [20] Ms A, Songnian Zhou, Xiaohu Zheng, Jingwen Wang, and Pierre Delisle. Utopia: A load sharing facility for large, heterogeneous distributed computer systems. *Software: Practice and Experience*, 23, 12 1993.
- [21] O. Tatebe, Y. Morita, S. Matsuoaka, N. Soda, and S. Sekiguchi. Grid datafarm architecture for petascale data intensive computing. In 2nd

IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID'02), pages 102–102, 2002.

- [22] James E Kelley Jr. Critical-path planning and scheduling: Mathematical basis. *Operations research*, 9(3):296–320, 1961.
- [23] Fabrizio Marozzo, Domenico Talia, and Paolo Trunfio. A workflow management system for scalable data mining on clouds. *IEEE Transactions On Services Computing*, 11(3):480–492, 2018.
- [24] Francisco Rodrigo Duro, Javier Garcia Blas, and Jesus Carretero. A hierarchical parallel storage system based on distributed memory for large scale systems. In *Proceedings of the 20th European MPI Users' Group Meeting, EuroMPI '13*, page 139–140, New York, NY, USA, 2013. Association for Computing Machinery.
- [25] Fabrizio Marozzo, Domenico Talia, and Paolo Trunfio. Js4cloud: Script-based workflow programming for scalable data analysis on cloud platforms. *Concurrency and Computation: Practice and Experience*, 27, 06 2015.
- [26] Robert D Blumofe and Charles E Leiserson. Scheduling multithreaded computations by work stealing. *Journal of the ACM (JACM)*, 46(5):720–748, 1999.
- [27] Tonglin Li, Xiaobing Zhou, Kevin Brandstatter, Dongfang Zhao, Ke Wang, Anupam Rajendran, Zhao Zhang, and Ioan Raicu. Zht: A light-weight reliable persistent dynamic scalable zero-hop distributed hash table. 05 2013.
- [28] Iman Sadooghi, Ke Wang, Dharmit Patel, Dongfang Zhao, Tonglin Li, Shiva Srivastava, and Ioan Raicu. Fabriq: Leveraging distributed hash tables towards distributed publish-subscribe message queues. 12 2015.
- [29] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
- [30] Weiwei Chen and Ewa Deelman. Workflowsim: A toolkit for simulating scientific workflows in distributed environments. pages 1–8, 10 2012.
- [31] Shishir Bharathi, Ann Chervenak, Ewa Deelman, Gaurang Mehta, Mei-Hui Su, and Karan Vahi. Characterization of scientific workflows. pages 1 – 10, 12 2008.



SALVATORE GIAMPÀ is a research fellow of computer engineering at the University of Calabria, Italy. He received a Master's degree in computer engineering in 2019. His research interests include distributed and parallel computing, programming framework and software engineering.



LORIS BELCASTRO is a research fellow of computer engineering at the University of Calabria, Italy. He received a Ph.D. in Information and Communication Engineering at the University of Calabria. In 2012 he held a scholarship at the Institute of High-Performance Computing and Networking of the Italian National Research Council (ICAR-CNR). His research interests include cloud computing, social media and Big Data analysis, distributed knowledge discovery, and data mining.



FABRIZIO MAROZZO is an assistant professor of computer engineering at the University of Calabria. He received a Ph.D. in Systems and Computer Engineering at the University of Calabria. In 2011-2012 he visited the Barcelona SuperComputing Center for a research internship with the Grid Computer Research group in Computer Sciences department. He is serving as associate editor the IEEE Access and the International Journal of Intelligent Systems Technologies and Applications.

His research interests include distributed systems, data mining, cloud computing, social media and Big Data analysis, and peer-to-peer networks.



DOMENICO TALIA is a professor of computer engineering at the University of Calabria. He is a member of the editorial boards of *Future Generation Computer Systems*, *IEEE Transactions on Parallel and Distributed Systems*, the *International Journal of Web and Grid Services*, the *Journal of Cloud Computing—Advances, Systems and Applications*, *Scalable Computing: Practice and Experience*, the *International Journal of Next-Generation Computing*. His research interests include parallel and distributed data mining algorithms, cloud computing, grid services, distributed knowledge discovery, peer-to-peer systems, and parallel programming models.



PAOLO TRUNFIO is an associate professor of computer engineering at the University of Calabria. In 2007 he was a visiting researcher at the Swedish Institute of Computer Science (SICS) in Stockholm. In 2001-2002 he was a research collaborator at the Institute of Systems and Computer Science of the Italian National Research Council (ISI-CNR). He is in the editorial board of *Future Generation Computer Systems*, *IEEE Transactions on Cloud Computing*, and *Journal of Big Data*. His

research interests include cloud computing, social-media analysis, service-oriented architectures, distributed knowledge discovery, and peer-to-peer systems.