



Resource provisioning for data-intensive applications with deadline constraints on hybrid clouds using Aneka

Adel Nadjaran Toosi^{*}, Richard O. Sinnott, Rajkumar Buyya

The Cloud Computing and Distributed Systems (CLOUDS) Laboratory, School of Computing and Information Systems, The University of Melbourne, Australia



HIGHLIGHTS

- A new data-aware provisioning algorithm is proposed to meet user-defined deadline requirements for data-intensive applications. The proposed algorithm takes into account available bandwidth and data transfer time.
- The proposed provisioning algorithm is integrated into the Aneka platform. Aneka is extended to support the Microsoft Azure Resource Manager (ARM) deployment service model.
- In an actual hybrid cloud environment, we evaluate the proposed algorithm's ability in meeting deadlines for a case study data-intensive application in smart cities context.

ARTICLE INFO

Article history:

Received 3 February 2017
Received in revised form 26 April 2017
Accepted 30 May 2017
Available online 24 June 2017

Keywords:

Dynamic provisioning
Hybrid cloud
Aneka cloud application platform
Deadline-driven scheduling
Data locality
Network bandwidth
Data-intensive applications

ABSTRACT

Cloud computing has emerged as a mainstream paradigm for hosting various types of applications by supporting easy-to-use computing services. Among the many different forms of cloud computing, hybrid clouds, which mix on-premises private cloud and third-party public cloud services to deploy applications, have gained broad acceptance. They are particularly relevant for applications requiring large volumes of computing power exceeding the computational capacity within the premises of a single organization. However, the use of hybrid clouds introduces the challenge of how much and when public cloud resources should be added to the pool of resources – and especially when it is necessary to support quality of service requirements of applications with deadline constraints. These resource provisioning decisions are far from trivial if scheduling involves data-intensive applications using voluminous amounts of data. Issues such as the impact of network latency, bandwidth constraints, and location of data must be taken into account in order to minimize the execution cost while meeting the deadline for such applications. In this paper, we propose a new resource provisioning algorithm to support the deadline requirements of data-intensive applications in hybrid cloud environments. To evaluate our proposed algorithm, we implement it in Aneka, a platform for developing scalable applications on the Cloud. Experimental results using a real case study executing a data-intensive application to measure the walkability index on a hybrid cloud platform consisting of dynamic resources from the Microsoft Azure cloud show that our proposed provisioning algorithm is able to more efficiently allocate resources compared to existing methods.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Data-intensive applications involving the analysis of large datasets have become increasingly important as many areas of science and business are facing thousand-fold increases in data volumes [1]. The explosive growth of data is mainly driven by the rapid expansion of the Internet, smart cities, social networks,

e-commerce, and widespread usage of high-throughput instruments, sensor networks, Internet of Things (IoT) devices, accelerators, and supercomputers. This expansion forms a voluminous amount of structured and unstructured data, known as big data, that needs to be processed to be useful [1]. The ability to analyze and process such large quantities of data has become an important and challenging mission for many fields.

Cloud computing [2] platforms are becoming one of the most preferred ways of hosting data-intensive applications. Challenges posed by big data can be overcome with the aid of cloud computing services offering the illusion of an infinite pool of highly reliable, scalable, and flexible computing, storage, and network resources.

^{*} Corresponding author.

E-mail address: adel.nadjaran@unimelb.edu.au (A. Nadjaran Toosi).

However, in many cases, data is available in local IT infrastructure with limited processing capacity, for example, a small cluster or resources from local area networks (desktop grids). Therefore, it is not time or cost effective to transfer the whole dataset to clouds to be processed. To tackle this issue, the *cloud bursting* model can be used in which an application runs in a private infrastructure and bursts onto a public cloud when more resources are required. This model has found broad acceptance due to its benefits such as cost reduction and dealing with issues related to the location of sensitive data [3].

To achieve the vision of cloud bursting, hybrid cloud *middleware* is required to acquire and release resources from both local infrastructures and external cloud providers in a seamless fashion [4]. It is essential for such hybrid cloud middleware to make efficient decisions regarding the workloads that must be outsourced to the public cloud based on the timing and number of externally provisioned resources to meet deadline constraints of applications. In such a setting, however, building a middleware that jointly minimizes cost and meets the deadline for applications is far from trivial [5].

There is a large body of literature aimed at cost and execution time minimization of running computational tasks in hybrid cloud environments. These studies mostly overlook aspects such as data locality, the impact of network bandwidth constraints, and data transfer time which significantly affect the time and cost performance of the scheduling. This is exacerbated for data-intensive applications where the data transfer time to the external cloud is often comparable to the computational time. In this paper, one of our main goals is to take these aspects into consideration for scheduling and resource provisioning of deadline-driven data-intensive applications in hybrid cloud environments.

In this context, *Platform-as-a-Service* (PaaS) solutions offer various tools to implement scheduling and resource provisioning policies in hybrid clouds. We exploit the Aneka platform [6] to implement the proposed solution in this paper. Aneka is a PaaS solution providing a middleware for the development and deployment of applications in hybrid and multi-clouds. Aneka provides application developers with Application Programming Interfaces (APIs) for transparently harnessing and exploiting the physical and virtual computing resources in heterogeneous networks of workstations, clusters, servers, and data centers. Earlier version of Aneka had many features supporting multi-cloud and hybrid computing. However, to provide wider support for scheduling and resource provisioning of data-intensive applications, we incorporate additional functionalities into Aneka.

In this paper, we make the following **key contributions**:

- We propose a new data-aware provisioning algorithm meeting the deadline requirements of applications executing in hybrid cloud environments. The proposed algorithm makes provisioning decisions by taking into account the data transfer time of scheduling tasks onto public cloud resources. This significantly affects data-intensive applications requiring large amount of data transfer. The main novelty of our approach is that while the solution takes into account bandwidth and data locality, it continually and dynamically updates scaling decisions based on the changes in the average runtime of the tasks and data transfer rates.
- The proposed algorithm is plugged into the Aneka platform that allows dynamically adding and removing resources from public clouds into the Aneka resource pool to meet user-defined application deadline requirements.
- Aneka is extended to support dynamic resource provisioning capabilities based on the Microsoft Azure Resource Manager (ARM) deployment service model.

- In an actual hybrid cloud environment built using local resources (desktop machines) and Azure virtual machines, we compared our method with existing approaches and demonstrated Aneka and its new provisioning algorithm's ability to meet deadlines for data-intensive applications. As a case study, a data-intensive application in the smart cities context is employed to measure the walkability index for different neighborhoods of the city of Melbourne utilizing spatial analysis of a large dataset [7].

The rest of the paper is organized as follows: Section 2 presents the motivation for this work and defines the problem domain. Section 3 outlines a general overview of the Aneka framework and describes the dynamic provisioning mechanisms of Aneka. Our proposed algorithm for deadline-driven data-aware resource provisioning is described in Section 4 and its realization in Aneka is discussed in Section 5. Section 6 is dedicated to a performance evaluation of the proposed algorithm. It discusses the hybrid cloud testbed built on top of computing resources from desktop grids and the Microsoft Azure cloud. Then, it describes the data-intensive case study application focused on measuring the walkability index and the assorted experimental results. Section 7 presents related work. Finally, Section 8 presents conclusions and offers future directions.

2. Motivation and problem domain

One of the main challenges for efficient scaling of applications is the location of the data relative to the available computational resources [8]. Co-locating data and computation is evidently ideal in terms of performance especially for data-intensive applications. However, this is not always feasible for various reasons. For example, data might be located in the storage nodes of the user's local organizational infrastructure (e.g., a cluster or desktop grid) with limited or overloaded computational resources and the user facing deadline constraints may prefer to leverage on-demand computing resources from a public cloud provider to reduce the execution time of the application.

In the above particular scenario, it may not be ideal for the user to move the entire dataset to the cloud as the data transfer time, due to the data size and network bandwidth, might dominate over the performance gain resulting from utilizing external CPUs. Moving data to distant computational resources, in particular for big data and data-intensive applications, to get access to more CPUs is often inefficient and can become the bottleneck in many cases. Therefore, any scheduling and resource provisioning algorithm aimed at improving data-intensive application performance by dynamic acquisition of cloud resources must take into account the time and amount of data movement. In other words, attempts to address the scheduling problem of data-intensive application would not be successful if they take into account computation separately from data movement. Accordingly, we focus on data-aware scheduling of data-intensive application considering data locality as well as monetary and performance costs of transferring data that has been neglected by many other scheduling methods in the literature.

Specifically, we focus on characteristics and requirements of hybrid cloud schedulers for executing deadline-constrained Bag-of-Tasks applications having large volumes of data. We assume that the application workload consists of a number of parallel tasks that each can run on an independent computing node. In addition, each task is associated with a dataset residing within the local infrastructure of the user that has to be fully transferred to the public cloud for those tasks to be executed on externally provisioned computational resources. Finally, the application has a deadline by which all its tasks must finish their execution.




Fig. 1. Aneka programming models.

In a hybrid cloud setting, the execution of an application happens through cloud bursting deployment models. Cloud bursting allows an application to run in a private data center and burst into a public cloud when more resources are required to meet a given deadline. In this paper, the application is primarily scheduled on private resources allocated from organizational infrastructure based on a best-effort algorithm. The scheduling algorithm needs to compute the time left for the deadline based on the average runtime of tasks. Extra resources from public clouds are dynamically allocated if the scheduling algorithm determines the number of resources (locally) acquired by the application is insufficient to meet the deadline. This process must take place repeatedly to continually update the average runtime of tasks. Note that in the remaining part of the paper, we use the term *runtime* to refer to the time period required to execute a task and the term *execution time* to refer to the total execution time of the application.

The main new added feature compared to our previous work is that the scheduler explicitly takes into account the size and transfer time of input/output data for the estimation of the required resources. We employ the Aneka middleware as the basis for supporting the proposed scheduling and provisioning algorithms to transparently execute the application in a hybrid cloud setting.

3. Aneka and dynamic resource provisioning

Aneka [6] is a software platform and framework facilitating the development and deployment of distributed applications onto clouds. It offers a collection of tools to build, control, and monitor cloud environment. The Aneka cloud built up this way can be composed of a collection of heterogeneous resources from a public cloud virtual infrastructure available through the Internet, a network of computing nodes in the premises of an enterprise, or a combination of both. Aneka provides developers with Application Programming Interfaces (APIs) for transparently exploiting physical and virtual resources in the Aneka cloud. Developers express the logic of applications using programming models and define runtime environments on top of which applications are deployed and executed. As shown in Fig. 1, Aneka currently supports four different programming models [6]: *Bag of tasks model*, *Distributed threads model*, *MapReduce model*, and *Parameter sweep model*.

3.1. Aneka architecture

The core components of the Aneka framework are designed and implemented in a service-oriented fashion. We briefly describe the architecture and the fundamental services that comprise the Aneka platform. Following this, we focus on the scheduler and provisioning services that are central to this paper.

Fig. 2 provides a layered view of the Aneka components. Aneka provides a runtime environment for executing applications by leveraging heterogeneous resources on the underlying *infrastructure* built on the top of computing nodes employed from network of workstations, clusters, grids, and data centers. In other words, the

infrastructure layer is a collection of nodes hosting components of Aneka middleware.

The *middleware* provides a collection of services for interactions with the Aneka cloud. The container represents the unit of deployment of Aneka clouds and the runtime environment for services. The core functionalities residing in the *Platform Abstraction Layer (PAL)* constitute the basic services that are used to control the infrastructure of Aneka clouds. It provides a uniform interface for management and configuration of nodes and the containers instances deployed on them in the infrastructure layer. Middleware is composed of two major components representing the building blocks of Aneka clouds: the *Aneka Daemon* and *Aneka Container*. Each node hosts the Aneka daemon and one or more Aneka container instances. The daemon is a management component controlling the container instances installed on the particular node. A node forms the infrastructure layer running the Aneka master container which plays the role of resource manager and application scheduler. Nodes running Aneka worker containers are responsible for processing and executing work units of the applications. In addition, each container provides a messaging channel for accessing features of different services provided by the container. There are three classes of services characterizing the container:

1. *Execution services*: are responsible for scheduling and executing applications. Specialized implementations of these services are defined for execution of work units of each programming model supported by Aneka.
2. *Foundation services*: are in-charge of metering applications, allocating resources, managing the collection of available nodes, and keeping the services registry updated.
3. *Fabric services*: provide access to the physical and virtualized resources managed by the Aneka cloud. The *Resource Provisioning Service (RPS)* enables horizontal scaling out and allows for elastic and dynamic growth and shrinkage of the Aneka cloud to meet Quality of Service (QoS) requirements of applications.

The services of the middleware are accessible through a set of interfaces and tools in the *development and management* layer. The *Software Development Kit (SDK)* embodies a collection of abstractions and APIs for definition of applications and leveraging existing programming models. The *Management Kit* contains a collection of tools for management, monitoring, and administration of Aneka clouds. All the management functions of the Aneka cloud are made accessible through the *Management Studio*, a comprehensive graphical environment providing a global view of the cloud for administrators.

3.2. Aneka scheduling and dynamic resource provisioning

Dynamic provisioning is the ability to dynamically acquire resources and integrate them into existing infrastructures and software systems. In the most common case, resources are *Virtual Machines (VMs)* acquired from an *Infrastructure-as-a-Service (IaaS)* cloud provider. Dynamic provisioning in Aneka happens as part of the Fabric Services by offering provisioning services for allocating virtual nodes from public cloud providers to complement local resources. This is mainly achieved as a result of the interaction between two services: the *Scheduling Service* and the *Resource Provisioning Service*. The former triggers on-demand provisioning requests based on the system status and the requirements of applications, while the latter is responsible for interacting with IaaS providers to instantiate VMs and deploy Aneka containers to meet the requests.

Execution of applications in Aneka happens through allocating tasks to the available set of resources in a dynamic fashion using the existing scheduling algorithms. Scheduling algorithms might




Fig. 2. Aneka framework overview [6].

be designed to leverage dynamic provisioning to cope with the application or system requirements. The scheduling algorithm makes decisions regarding when and how many resource allocations must take place to meet the application QoS requirements.

Aneka supports interactions with different resource providers, e.g., Amazon Elastic Computer Cloud (EC2), Microsoft Azure, XenServer, and GoGrid, using its dedicated provider-specific *resource pool* component. The main operations performed by this component are the translation of provisioning requests into provider specific requests, controlling the life cycle of VMs, and shutting them down when they are no longer needed. The life cycle of resource pools and redirecting provisioning requests, their release, or directing queries to the appropriate pool is the responsibility of the *pool manager* component. The pool manager also notifies the provisioning service when a dynamic resource is activated and terminated. Fig. 3 illustrates a schematic overview of Aneka's dynamic provisioning mechanism.

Aneka features several provisioning algorithms that are designed to support dynamic provisioning of virtual resources. Among these, the algorithms proposed in [4,6] are designed to leverage dynamic resources to meet the deadline requirements of Bag-of-Tasks applications. The algorithm proposed by Veccholia et al. [6], which we call *Default* here, makes an estimation of the expected completion time of the application with currently available resources and if the expected completion time is later than the deadline defined in the Quality of Service parameters of the application, it requests extra resources from the public cloud to complete the application within given deadlines. This algorithm provides a best effort strategy for meeting the required deadlines based on the average task runtime estimation. Since the *Default* algorithm ignores the deployment time of resources (e.g., VM startup time) during the calculation of number of extra required resources,

Calheiros et al. [4] proposed an improved provisioning algorithm, which we call *Enhanced* here, that dynamically leases resources to meet deadlines while it takes deployment time of resources into account. The *Enhanced* algorithm is designed to utilize Amazon EC2 Spot Instance resources with on average higher deployment time but lower budget than the *Default* algorithm. One of the key differences between our proposed provisioning algorithm in the next Section and these existing algorithms is that we consider the network bandwidth and data transfer time during the calculation of the required extra resources. In Section 6, we compare our proposed algorithm with the *Default* and *Enhanced* algorithms based on different parameters such as application execution time, the number of resources launched on the public cloud, and the cost of resources.

4. Deadline-driven data-aware resource provisioning algorithm

Suppose that a private cloud with a limited number of resources in the local infrastructure is available for execution of a data-intensive Bag-of-Tasks application. Since the number of tasks that can be running concurrently on the private cloud is limited, to meet the deadline requirement of the application, extra resources from the public cloud need to be acquired to scale out available resources. We assume that the public cloud provider is able to fulfill all requests and thus has an infinite number of resources available from the user's perspective. Note that the network *bandwidth* available between the private and public cloud is limited and can impose a significant amount of data transfer time for each task running on the public cloud resource.

We assume that the application's workload consists of a number of trivially parallel tasks, each requiring specific input data files




Fig. 3. A schematic overview of Aneka's dynamic provisioning.

Algorithm 1 Data-aware Provisioning Algorithm.

```

1: privateCores ← private cores available for the application;
2: avgTaskRuntime ← Average task runtime on a private core;
3: timeRemaining ← Time to application deadline;
4: totalTasks ← Total number of tasks in the application;
5: tasksCompleted ← Total number of tasks completed so far;
6: startupTime ← Startup time of a resource (VM);
7: tasksInPrivate ←  $\lfloor \frac{\text{timeRemaining}}{\text{avgTaskRuntime}} \times \text{privateCores} \rfloor$ ;
8: tasksRemaining = (totalTasks – tasksCompleted – tasksInPrivate)+;
9: totalTransferTime ← tasksRemaining ×  $\frac{\text{taskInputDataSize}}{\text{upBandwidth}}$ ;
10: actualTimeRemaining ← (timeRemaining – startupTime – totalTransferTime)+;

11: avgTaskRuntimeOnPublic ← Average task runtime on public core;
12: provisionedCores ← Current dynamically provisioned cores;
13: totalExecutionTime ← tasksRemaining × avgTaskRuntimeOnPublic;
14: tasksPerCore ←  $\lfloor \frac{\text{actualTimeRemaining}}{\text{avgTaskRuntimeOnPublic}} \rfloor$ ;
15: if tasksPerCore < 1 then
16:   if toGrow then
17:     totalCoresRequired ← provisionedCores;
18:   else
19:     totalCoresRequired ← provisionedCores – 1
20:   end if
21: else
22:   totalCoresRequired ←  $\lceil \frac{\text{totalExecutionTime}}{\text{tasksPerCore} \times \text{avgTaskRuntimeOnPublic}} \rceil$ ;
23: end if
24: extraResources ←  $\lceil \frac{\text{totalCoresRequired} - \text{provisionedCores}}{\text{numberOfCoresPerResource}} \rceil$ ;

```

located in the local infrastructure. Bag-of-Tasks applications with independent tasks are used in a variety of scenarios, especially when the same piece of computation logic must be executed over a large volume of data, e.g., Monte Carlo simulations, data mining algorithms, parameter sweep applications [9]. This makes them suitable for hybrid cloud scenarios as tasks running on two separate clouds do not need to communicate with each other.

Algorithm 1 presents the newly proposed dynamic resource provisioning approach for Aneka, called *Data-aware* provisioning.

The Data-aware algorithm takes into account the available bandwidth and the data size associated with each task and calculates the number of extra resources required to meet the deadline constraints of the application. The algorithm is executed when any of the following conditions are observed: (i) a task from the application is queued, and (ii) execution of a task completes. The variable toGrow at line 16 is True if the former condition happens and is False otherwise. If completion of a task triggers the algorithm and the computed number of extra resources by the algorithm is negative, a release request for the task allocated resource is submitted to the *Resource Provisioning Service*; otherwise, extra resources computed by the algorithm are requested to be added to the pool. Moreover, in order to reduce unnecessary calls of the algorithm, we only execute the algorithm in the growing mode, i.e., the average runtime of tasks is increased in comparison to the previous round of algorithm execution. Similarly, we call it in the shrinking mode if the average runtime of the task is decreased in comparison to the previous round of algorithm execution. For the sake of brevity, these conditions are not shown in the algorithm.

The Data-aware algorithm checks if the currently available resources are sufficient for the completion of the application tasks within the given deadline based on estimation of the average runtime of tasks on the private resources (avgTaskRuntime). Note that avgTaskRuntime includes the data transfer time in the calculation as we assume that the data transfer time within the private cloud is insignificant compared to the task execution time. Therefore, we do not consider a separate variable to capture that in the algorithm. The algorithm first updates timeRemaining based on the left time to the deadline. Then it computes the number of tasks that can be completed on the private resources within the left time to the deadline (Line 4). Then, it calculates the number of remaining tasks (Line 8). Here (x)⁺ means max(0, x). These remaining tasks must be scheduled on the dynamic resources. Since the execution of tasks on resources from the public cloud requires transferring data from the local storage to the allocated VMs, the

algorithm first calculates the total transfer time for the tasks that should be executed on the public cloud resources (Line 9). In Line 10, the actual remaining time which can be effectively used for the execution of the tasks on dynamic resources is computed by subtracting the total transfer time and the start-up time of resources from the remaining time to the deadline. Note that `avgTaskRuntimeOnPublic` in Line 11, contrary to `avgTaskRuntime`, is only calculated based on the actual time tasks are being executed on the public cloud resources and does not include any associated data transfer time, i.e., it is the time period from when the task starts execution on the compute node (after all input data is available) up to the moment its execution is over. In line 13, the total time required to execute all remaining task on a single public CPU core is calculated and stored in the `totalExecutionTime` variable.

By dividing the actual time remaining by the average runtime of tasks on public cloud resources, the number of tasks each CPU core on the public resources can execute is calculated. If the number of tasks can be executed on each core in the public resources is at least one, then the total number of required cores can be estimated by dividing `totalExecutionTime` by the result of tasks per core multiplied by the average runtime of the task (Line 22). Otherwise, there is not enough time for allocating new resources and the algorithm sets the number of required CPU cores to that of the already provisioned cores in the pool (Line 17). If this is a shrinking mode, the task allocated cores will be nominated for release (Line 19). Line 24, calculates extra resources that must be added to or removed from the pool by the ratio of extra required cores to the number of cores per resource.

5. Realization of data-aware provisioning algorithm in Aneka

To support execution of applications in hybrid clouds, it is important for Aneka to have access to public cloud resources. The Aneka resource provisioning service currently supports provisioning for providers such as Amazon EC2, GoGrid, and Microsoft Azure. Microsoft Azure has undergone a significant transformation in recent years and as a result, there are two different sets of APIs for resource management and deployment in Azure: Azure Resource Manager (ARM) and Classic. Aneka originally supported the Classic deployment model in which each resource (e.g., storage disk, VM, Public IP address, etc.) existed independently and there was no way to group related resources together. In 2014, Azure introduced ARM, which added the concept of a resource group as a container for resources that share a common lifecycle. As part of implementing Data-aware provisioning algorithm, we added an Azure-specific resource pool based on ARM APIs into Aneka which handles provisioning requests for Azure.

Since in the hybrid cloud environment, computational nodes are scattered throughout multiple networks (i.e., private and public cloud networks), public IP addresses are required to make the communication between nodes possible. Providing public IP addresses for all nodes, in particular for private resources in the organizational infrastructure, is not always feasible. Therefore, we used Azure Virtual Private Network (VPN) to solve this issue. We configured a Point-to-Site (P2S) Virtual Network in Azure which allows a secure connection from an individual client computer to the virtual network. Point-to-Site is a VPN connection over SSTP (Secure Socket Tunneling Protocol) and does not require a VPN device or a public-facing IP address to work. Fig. 4 shows a sample Azure Point-to-Site virtual network configuration.

Aneka provides interfaces for installation of Aneka containers on static resources (i.e., private cloud resources) via its Aneka Management Studio. However, in dynamic resource provisioning when machines are added on-demand by instantiating VMs, the installation must be done automatically without any human intervention. Therefore, we create virtual machine images on Azure

for Aneka Workers with an initial configuration for containers. To complete the installation process, instantiated VMs are further configured using PowerShell scripts to set the Master container IP and Port, configure IP addresses in configuration files and to start the container service.

Finally, to support execution of data-intensive applications in hybrid cloud environments, we incorporate the Data-aware algorithm in Aneka. Aneka offers abstract interfaces that can be used to implement new resource provisioning algorithms. Our algorithm is implemented as a new scheduling algorithm by implementing the `ISchedulingAlgorithm` interface. The algorithm invokes the provisioning service to add extra resources to the pool of available resources based on the QoS requirement of the application.

6. Performance evaluation

6.1. Hybrid cloud setup

The experimental testbed used to evaluate the performance of the proposed resource provisioning and scheduling algorithms is a hybrid cloud environment constituting of two desktop machines (one master and one slave) residing at The University of Melbourne and dynamic resources provisioned from Microsoft Azure. A schematic view of the hybrid cloud in which the experiments are carried out is depicted in Fig. 5 and the configurations of machines used in the experiment are shown in Table 1.

Public cloud resources are dynamically provisioned from Microsoft Azure cloud when local resources are not able to meet application deadlines. Public cloud resources are deployed on Azure Australia Southeast region using Standard DS1 VM instances with the configuration shown in Table 1. The initial estimated start-up time of Azure VMs used by the proposed algorithm is 250 s. However, this is updated by a new value each time a VM is added to the resource pool.

The master and worker machines are connected through a high bandwidth LAN connection that imposes a negligible data transfer time between the local machines. On the other hand, the data transfer from the data repository on the master node to dynamic resources on the cloud is performed over the Internet. Since, the Internet connection on the master is relatively high for the intended experiments; we used `NetLimiter`¹ Version 4 to limit both the up-link and down-link bandwidth of the master node to 2 MB/s.

6.2. Walkability application and dataset

The data-intensive application used for the experiments is a Bag-of-Tasks for measuring a walkability index [7]. A walkability index is used to assess how walkable a given neighborhood is based on factors such as road connectivity, gross dwelling density and the land use mix in the area. The information can be used for better understanding the growing urban challenges such as obesity and increasing dependence on cars. The walkability index application is part of the AURIN project which supports nationwide urban and built environment research across Australia [10].

In our experiments, we use the walkability application to provide walkability indexes for 220 different neighborhoods in the city of Melbourne. The walkability application suits the purpose of our experiments since it is data-intensive and it can be broken into independent tasks, each computing a walkability index for a neighborhood. The test application contains 55 tasks, each calculating walkability indexes for four different neighborhoods. The input data for each task contains geospatial datasets for road

¹ NetLimiter, <https://www.netlimiter.com/download>.




Fig. 4. A sample Azure Point-to-Site virtual network.




Fig. 5. Hybrid cloud testbed.

Table 1
Configuration of machines used in the experiments.

Machine	Type	CPU (GHz)	Cores	Memory (GB)	OS
Master	Intel Core i7-4790	3.60	8	16	Windows 7
Worker	Intel Core i7-2600	3.40	8	8	Windows 7
Azure Instances	Standard DS1	2.4	1	3.5	Windows Server 2012

networks provided by the *Public Sector Mapping Agent (PSMA) Australia*² and points representing the center of each neighborhood. The size of input data for each task is 130 MB including the Java executable JAR file. The input dataset must be transferred to the computing machine from the Aneka data repository residing in the master node. The output is the computed walkability score for each neighborhood and its size for each task is negligible. Fig. 6 displays a sample snapshot of output data for several neighborhoods visualized on the map according to the computed walkability indexes. Each area is colored based on the walkability score with blue marked area noted as the most walkable and red representing the least walkable.

6.3. Experimental results

In order to evaluate the performance of the proposed data-aware algorithm, we first submit the walkability application to

Aneka for execution without setting a deadline for the application. Later, we repeat the execution with different deadlines, showing how the proposed algorithm behaves. All experiments are repeated for the Default [6] and Enhanced [4] algorithms for the sake of comparison. By running the preliminary experiment, we are able to estimate the expected execution time of the application, which allows us to impose a deadline triggering the resource provisioning in other experiments. We identified that the execution time of the application without setting a deadline which only uses private (local) resources takes 45.4 min. Therefore, we considered test scenarios with four different deadline values of 35, 40, 45 and 50 min. Note that the deadline value of 35 min is the tightest deadline we trial in our experiments since all algorithms failed to meet the deadline within 30 min.

Fig. 7 and Table 2 show the results of the application execution under different deadlines in two different forms. As shown by Fig. 7(a), the proposed data-aware algorithm meets the deadline in all scenarios. However, the Default and Enhanced algorithms miss the deadline under tight deadline constraints (Bars marked

² PSMA, <https://www.pdma.com.au/>.




Fig. 6. A snapshot of visualized Walkability indexes computed for a set of neighborhoods around Melbourne. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

by “x”), i.e., both algorithms violate the deadline constraint when the deadline is set to 35 min whilst the Default algorithm misses the deadline when it is set to 40 min. The key reason is that these algorithms rely on only a single variable for measuring average runtime of tasks (the value includes the data transfer time) to allocate dynamic resources that is significantly different in our scenario for private and public cloud resources. This difference is due to the dissimilarity of the data transfer time for the tasks executed on the private and public cloud resources which is dependent on the difference in the available bandwidth for each case. We conclude that our new algorithm is able to meet strict application deadlines by taking into account the start-up time of the VMs and data transfer time, while the Default and Enhanced algorithms fail to do the same as they underestimate the overall runtime by neglecting significantly different intra and inter cloud data transfer time in the hybrid cloud setting.

The experimental results also show that Data-aware algorithm instantiates a lower number of VMs in comparison to the two other algorithms when deemed necessary (e.g., at deadline values of 40, 45, and 50 min). This can be justified by the fact that the Data-aware algorithm takes into account the data transfer time and is less affected by the fluctuation of the average runtime of tasks. The fluctuation of the average task runtime leads to more hasty decisions by the Default and Enhanced algorithms to terminate VMs that may subsequently need to be instantiated again shortly after the termination. It is noted that the only case that the Data-aware algorithm instantiates a higher number of VMs than other algorithms is when the deadline is 35 min (i.e., 19 VMs compared to 10 and 11 for Default and Enhanced, respectively). However, as noted, both other competing algorithms fail to meet the deadline in this case.

As shown by Fig. 7(a) and (c) and Table 2, the proposed Data-aware algorithm not only guarantees the deadline but also optimizes the (monetary) cost spent on allocation of public cloud resources from Azure. Since Azure VMs are billed on a *per-minute* basis, we report the total number of minutes that dynamically added resources are kept running in the resource pool by each algorithm as an indication of cost. Compared to Data-aware, the Default and Enhanced algorithms run dynamically provisioned VM resources for a higher number of minutes in all cases where

they meet the deadline. This happens because of both the higher expenditure on total start-up time of VMs due to more frequent instantiation and termination of VMs and the over-provisioning of resources due to late or wrong decisions regarding the addition of dynamic resources.

As stated before, the estimated execution time of the application under no deadline constraints is roughly 45 min. Therefore, we expect that none of the algorithms provision any dynamic resources when the deadline is set to a value higher than 45 min. However, all algorithms witness some degree of over-provisioning when the deadline is set to 50 min or even 45 min. The reason behind this is the primary estimate provided for the tasks runtime which is higher than the actual average runtime in our case can lead to wrong initial decisions made by algorithms and adding dynamic resources to the pool of resources. In the Aneka platform, whenever a resource is allocated to a task the mapping is irreversible unless a task or resource failure occurs. As a result, all algorithms use dynamically added VMs for the execution of at least one task if they add dynamic resources. As can be seen in both Fig. 7 and Table 2, the Data-aware algorithm incurs a lower amount of over-provisioning compared to other algorithms in case of over-estimation of task runtime in the initial setup of the scheduling algorithms.

Our conclusion is that providing Aneka scheduler with more accurate initial estimates of the runtime of tasks helps to minimize these over-provisioning issues. Moreover, considering the start-up of VMs that are particular to different infrastructures together with the size of the transferred data and network bandwidth are necessary in providing an efficient scheduling in the Aneka platform. This is particularly acute for data-intensive applications. In line with this, our experimental results show that our proposed algorithm is able to meet strict application deadlines with minimal budget expenditure by taking into account such factors.

7. Related work

Resource provisioning is one of the most challenging problems in the cloud environment [11]. Resources must be allocated dynamically according to the Quality of Service (QoS) requirements and workload changes of the application. Autonomic systems provide solution to this problem by offering the environment in which

Table 2
Experimental results for Default, Enhanced, and Data-aware algorithms considering different application deadlines.

Deadline (min)	Execution time (min)			Launched Azure VMs (#)			Total running time of VMs (min)		
	Default	Enhanced	Data-aware	Default	Enhanced	Data-aware	Default	Enhanced	Data-aware
35	44	42	35	10	11	19	182	133	513
40	40	37	38	8	10	7	105	146	135
45	40	44	44	8	10	3	86	104	61
50	41	42	43	7	7	1	82	104	32
$+\infty$	45	45	45	0	0	0	0	0	0




Fig. 7. (a) Execution time. (b) Number of launched Azure VMs and (c) Total running time of VMs for Default, Enhanced, and Data-aware algorithms considering different application deadlines. The X symbol shows a violated deadline.

dynamic resource provisioning for applications can be performed without human intervention [12]. To achieve the goal of autonomous systems, the *Monitor-Analyze-Plan-Execute* (MAPE) reference model proposed by IBM is of the most popular architectures. Ghobaei-Arani et al. [13] proposed a generic resource provisioning framework for cloud applications based on the MAPE architecture. In contrast to our approach, they focus on a single cloud application and do not investigate the impact of network latency, bandwidth constraints, and location of data on the dynamic resource provisioning.

The idea of using public cloud resources to expand the capacity of local infrastructure has been explored by many studies. Mateescu et al. [14] propose an architecture that provides a platform for the execution of High-Performance Computing (HPC) scientific applications. The cornerstone of the proposed architecture is the *Elastic Cluster* which makes an expandable hybrid cloud environment. Their approach differs from ours since we focus specifically on data-intensive applications and take into account the impact of data transfer times. Assunção et al. [15] analyze the trade-off between performance and usage costs of different provisioning algorithms for use of resources from the cloud to expand a cluster capacity. Similar to [14], they neglect the impact of data transfer time. Javadi et al. [16] propose failure-aware resource provisioning policies for hybrid cloud environments which they evaluated using a model-based simulation as opposed to our real case study performance evaluation. Xu and Zhao [3] propose a privacy-aware hybrid cloud framework which supports a tagging mechanism for the location of sensitive data. While they focus on compliance with the location of sensitive data, we focus on the data locality and data transfer time to compute the number of public cloud resources required by the application in order to meet deadlines. Belgacem and Chopard [17] conduct an experimental study of running a large, tightly coupled, distributed application over a hybrid cloud consisting of resources from Amazon EC2 clusters and an existing HPC infrastructure. They evaluated the overhead of using public cloud resources for a tightly coupled, massively parallel MPIs application in which tasks are communicating with each other without considering QoS related factors such as deadlines. However, in this work, we look into the design of dynamic resource provisioning algorithms meeting the deadline constraint of the locally coupled Bag-of-Tasks applications. Mattess et al. [18] present a provisioning algorithm for extending cluster capacity with Amazon EC2 *Spot Instances*. In contrast to our work, they focus on compute-intensive applications without considering network related costs and delays. Yuan et al. [19] propose a profit maximization model for private cloud providers by utilizing the temporal variation of prices in hybrid cloud. While similar to many others, they assume the time and cost related to data and network are negligible. The majority of these works focus largely on theoretical aspects and evaluate their method through simulation, while we focus on practical aspects and execute our case study on a real hybrid cloud environment using the Aneka platform with a real application.

Scheduling and resource provisioning techniques in a hybrid cloud for data-intensive applications where the data transfer time is comparable to computational time, adds new levels of complexities requiring addressing the impact of network latency, bandwidth constraints, as well as economic aspects such as costs and prices. A thorough survey on resource scheduling and provisioning in cloud environments has been conducted by Singh and Chana [20]. Bossche et al. [5] proposed scheduling algorithms to deal with cost optimization problem for deadline-constrained applications while taking into account data constraints, data locality and inaccuracies in task runtime estimates. Similar to our approach, their algorithm considers computational and data transfer costs as well as network bandwidth constraints. However, their work is a simulation-based study and does not consider the dynamic nature of networks and fluctuations in the runtime of tasks.

Kailasem et al. [21] propose a Hadoop-based cloud bursting framework for data-intensive workloads. They focus on checkpointing schemes to overlap data download with processing and upload. They consider the performance optimization of interactive jobs while we target Bag-of-Tasks jobs. They also ignore the uncertainty of job execution runtimes. Similar to our work, Bicer et al. [22] propose a resource allocation framework suitable for a hybrid cloud settings to support time and cost sensitive execution of data-intensive applications. Different from our work, their method is designed for Map-Reduce applications and is based on a feedback mechanism in which the compute nodes regularly report their performance. Malawski [23] proposed a mixed integer nonlinear programming model to tackle the problem of resource allocation on multiple heterogeneous clouds taking into account the cost of instances and data transfers. While we focus both on how many and when extra resources must be added in a dynamic fashion, they try to provide a single decision using a mixed integer non-linear programming model to tackle the resource provisioning problem in hybrid cloud environments. In contrast to our work, they have used simulation to evaluate their method.

Scheduling and resource provisioning in hybrid clouds has been researched for other types of application as well. Examples include big data analytics [24], workflows applications [25], online commerce [26], mobile phone applications [27] and compute intensive applications [28].

8. Conclusions and future directions

Supplementing on-premises private infrastructure of organizations with dynamically provisioned resources from public cloud providers introduces the problem of cost-efficiently executing applications. This paper presents a provisioning algorithm for scheduling deadline-constrained data-intensive applications while taking into account aspects such as data transfer time, the location of data, and the network bandwidth. This work builds upon previously proposed provisioning algorithms for the Aneka platform for developing and deploying scalable applications on the cloud. In this work, we propose a provisioning algorithm that computes the extra resources needed to complete application tasks within deadlines by considering aspects such as data locality, start-up time of public cloud resources, network bandwidth, and data transfer time. We demonstrate that our proposed algorithm is able to meet strict deadlines for a sample data-intensive application while minimizing cost and the total number launched instances compared to other existing algorithms. Contrary to other algorithms, the proposed algorithm measures the average runtime of tasks on public cloud resources as a separate variable and takes the data transfer time calculated based on the available bandwidth into account.

In this paper, we focused on scheduling and resource provisioning of Bag-of-Tasks applications with a set of trivially parallel tasks which can be executed independently of one another. However, there are other common types of data-intensive applications with a non-trivial workflow structure that includes precedence constraints and data dependencies between tasks. A meaningful future work is to extend our proposed algorithm for workflows with data dependencies. This is a complex problem and the major challenge in the design of such algorithms in addition to consideration of data transfer time and data locality is how to devise the task grouping and task assignment techniques that minimize inter-cloud communications.

Another future work in line with this contribution consists of developing provisioning policies that can support the integration of multiple clouds with different pricing and network latency in Aneka. We also plan to look into the scheduling of Bag-of-Tasks applications with tasks partly sharing input data and other QoS parameters such as budget constraints.

Future directions of data-intensive application scheduling include new algorithms for multi-cloud resource allocation, innovative provisioning algorithms honoring user requirements such as privacy and the location of sensitive data, energy efficient techniques for minimizing cost and carbon footprint. This can be further extended to support resource allocation techniques leveraging software-defined networks. This can also be explored in the context of various other programming paradigms such as Map Reduce model.

Acknowledgments

This work was partially supported by Australian Research Council (ARC) Future Fellowship (Grant no: FT120100545) and the Australia-India Strategic Research Fund (AISRF) (Grant no: AISRF08140). We thank Microsoft for providing access to the Azure IaaS infrastructure. We also thank William Voorsluys for his technical support for the execution of the case study application.

References

- [1] C.P. Chen, C.-Y. Zhang, Data-intensive applications, challenges, techniques and technologies: a survey on big data, *Inform. Sci.* 275 (2014) 314–347. <http://dx.doi.org/10.1016/j.ins.2014.01.015>.
- [2] R. Buyya, C.S. Yeo, S. Venugopal, J. Broberg, I. Brandic, Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility, *Future Gener. Comput. Syst.* 25 (6) (2009) 599–616. <http://dx.doi.org/10.1016/j.future.2008.12.001>.
- [3] X. Xu, X. Zhao, A framework for privacy-aware computing on hybrid clouds with mixed-sensitivity data in: *Proceedings of the IEEE International Symposium on Big Data Security on Cloud, 2015*, pp. 1344–1349 <http://doi.org/10.1109/HPCC-CSS-ICSS.2015.110>.
- [4] R.N. Calheiros, C. Vecchiola, D. Karunamoorthy, R. Buyya, The Aneka platform and QoS-driven resource provisioning for elastic applications on hybrid clouds, *Future Gener. Comput. Syst.* 28 (6) (2012) 861–870. <http://dx.doi.org/10.1016/j.future.2011.07.005>.
- [5] R.V. den Bossche, K. Vanmechelen, J. Broeckhove, Online cost-efficient scheduling of deadline-constrained workloads on hybrid clouds, *Future Gener. Comput. Syst.* 29 (4) (2013) 973–985. <http://dx.doi.org/10.1016/j.future.2012.12.012>.
- [6] C. Vecchiola, R.N. Calheiros, D. Karunamoorthy, R. Buyya, Deadline-driven provisioning of resources for scientific applications in hybrid clouds with aneka, *Future Gener. Comput. Syst.* 28 (1) (2012) 58–65. <http://dx.doi.org/10.1016/j.future.2011.05.008>.
- [7] R.O. Sinnott, W. Voorsluys, A scalable cloud-based system for data-intensive spatial analysis, *Int. J. Softw. Tools Technol. Trans.* 18 (6) (2016) 587–605. <http://dx.doi.org/10.1007/s10009-015-0398-6>.
- [8] I. Foster, Y. Zhao, I. Raicu, S. Lu, Cloud computing and grid computing 360-degree compared, in: *Proceedings of 2008 Grid Computing Environments Workshop, 2008*, pp. 1–10, ISSN: 2152-1085 <http://dx.doi.org/10.1109/GCE.2008.4738445>.
- [9] F.A. da Silva, H. Senger, Scalability limits of bag-of-tasks applications running on hierarchical platforms, *J. Parallel Distrib. Comput.* 71 (6) (2011) 788–801. <http://dx.doi.org/10.1016/j.jpdc.2011.01.002>. Special Issue on Cloud Computing.
- [10] R.O. Sinnott, C. Bayliss, A. Bromage, G. Galang, G. Grazioli, P. Greenwood, A. Macaulay, L. Morandini, G. Nogoarani, M. Nino-Ruiz, M. Tomko, C. Pettit, M. Sarwar, R. Stimson, W. Voorsluys, I. Widjaja, The Australia urban research gateway, *Concurr. Comput.: Pract. Exper.* 27 (2) (2015) 358–375. <http://dx.doi.org/10.1002/cpe.3282>. CPE-13-0325.R1.
- [11] M. Amiri, L. Mohammad-Khanli, Survey on prediction models of applications for resources provisioning in cloud, *J. Netw. Comput. Appl.* 82 (2017) 93–113. <http://dx.doi.org/10.1016/j.jnca.2017.01.016>.
- [12] S. Singh, I. Chana, R. Buyya, STAR: SLA-aware autonomic management of cloud resources, *IEEE Trans. Cloud Comput.* (2017). <http://dx.doi.org/10.1109/TCC.2017.2648788>.
- [13] M. Ghobaei-Arani, S. Jabbehdari, M.A. Pourmina, An autonomic resource provisioning approach for service-based cloud applications: a hybrid approach, *Future Gener. Comput. Syst.* (2017). <http://dx.doi.org/10.1016/j.future.2017.02.022>.
- [14] G. Mateescu, W. Gentsch, C.J. Ribbens, Hybrid computing—where HPC meets grid and cloud computing, *Future Gener. Comput. Syst.* 27 (5) (2011) 440–453. <http://dx.doi.org/10.1016/j.future.2010.11.003>.
- [15] M.D. de Assunção, A. di Costanzo, R. Buyya, A cost-benefit analysis of using cloud computing to extend the capacity of clusters, *Cluster Comput.* 13 (3) (2010) 335–347. <http://dx.doi.org/10.1007/s10586-010-0131-x>.

- [16] B. Javadi, J. Abawajy, R. Buyya, Failure-aware resource provisioning for hybrid cloud infrastructure, *J. Parallel Distrib. Comput.* 72 (10) (2012) 1318–1331. <http://dx.doi.org/10.1016/j.jpdc.2012.06.012>.
- [17] M.B. Belgacem, B. Chopard, A hybrid HPC/cloud distributed infrastructure: Coupling EC2 cloud resources with HPC clusters to run large tightly coupled multiscale applications, *Future Gener. Comput. Syst.* 42 (2015) 11–21. <http://dx.doi.org/10.1016/j.future.2014.08.003>.
- [18] M. Mattess, C. Vecchiola, R. Buyya, Managing peak loads by leasing cloud infrastructure services from a spot market, in: *Proceedings of the 12th IEEE International Conference on High Performance Computing and Communications HPCC, 2010*, pp. 180–188. <http://dx.doi.org/10.1109/HPCC.2010.77>.
- [19] H. Yuan, J. Bi, W. Tan, B.H. Li, Temporal task scheduling with constrained service delay for profit maximization in hybrid clouds, *IEEE Trans. Autom. Sci. Eng.* 14 (1) (2017) 337–348. <http://dx.doi.org/10.1109/TASE.2016.2526781>.
- [20] S. Singh, I. Chana, A survey on resource scheduling in cloud computing: issues and challenges, *J. Grid Comput.* 14 (2) (2016) 217–264. <http://dx.doi.org/10.1007/s10723-015-9359-2>.
- [21] S. Kailasam, P. Dhawalia, S.J. Balaji, G. Iyer, J. Dharanipragada, Extending MapReduce across Clouds with BStream, *IEEE Trans. Cloud Comput.* 2 (3) (2014) 362–376. <http://dx.doi.org/10.1109/TCC.2014.2316810>.
- [22] T. Bicer, D. Chiu, G. Agrawal, Time and cost sensitive data-intensive computing on hybrid clouds, in: *Proceedings of 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing CCGrid 2012, 2012*, pp. 636–643. <http://dx.doi.org/10.1109/CCGrid.2012.95>.
- [23] M. Malawski, K. Figiela, J. Nabrzyski, Cost minimization for computational applications on hybrid cloud infrastructures, *Future Gener. Comput. Syst.* 29 (7) (2013) 1786–1794. <http://dx.doi.org/10.1016/j.future.2013.01.004>.
- [24] F.J. Clemente-Castell, B. Nicolae, K. Katrinis, M.M. Rafique, R. Mayo, J.C. Fernandez, D. Loreti, Enabling big data analytics in the hybrid cloud using iterative mapreduce in: *Proceedings of the 8th IEEE/ACM International Conference on Utility and Cloud Computing UCC, pp. 290–299, 2015*. <http://dx.doi.org/10.1109/UCC.2015.47>.
- [25] M.-A. Vasile, F. Pop, R.-I. Tutueanu, V. Cristea, J. Koodziej, Resource-aware hybrid scheduling algorithm in heterogeneous distributed computing, *Future Gener. Comput. Syst.* 51 (2015) 61–71. <http://dx.doi.org/10.1016/j.future.2014.11.019>.
- [26] G. Lackermair, (2011) Hybrid cloud architectures for the online commerce in: *Procedia Computer Science, World Conference on Information Technology, Vol. 3, 2011*, pp. 550–555, ISSN1877-0509 <http://dx.doi.org/10.1016/j.procs.2010.12.091>.
- [27] H. Flores, S.N. Srirama, C. Paniagua, A generic middleware framework for handling process intensive hybrid cloud services from mobiles, in: *Proceedings of the 9th International Conference on Advances in Mobile Computing and Multimedia, MoMM '11, ACM, New York, NY, USA, 2011*, pp. 87–94. <http://dx.doi.org/10.1145/2095697.2095715>.
- [28] M. Brock, A. Goscinski, Execution of compute intensive applications on hybrid clouds, case study with mpiBLAST, in: *Proceedings of the Sixth International Conference on Complex, Intelligent, and Software Intensive Systems, 2012*, pp. 995–1000 <http://dx.doi.org/10.1109/CISIS.2012.109>.



Adel Nadjaran Toosi is Post-doctoral Research Fellow at the Cloud Computing and Distributed Systems (CLOUDS) Laboratory, School of Computing and Information Systems (CIS), the University of Melbourne, Australia. He received his B.Sc. degree in 2003 and his M.Sc. degree in 2006 both in Computer Science and Software Engineering from Ferdowsi University of Mashhad, Iran. He has done his Ph.D., supported by International Research Scholarship (MIRS) and Melbourne International Fee Remission Scholarship (MIFRS), at the CIS department of the University of Melbourne. Adel's thesis was nominated for CORE John

Makepeace Bennett Award for the Australasian Distinguished Doctoral Dissertation and John Melvin Memorial Scholarship for the Best Ph.D. thesis in Engineering. His current h-index is 14 based on the Google scholar Citations. His research interests include scheduling and resource provisioning mechanisms for distributed systems. Currently he is working on data-intensive application resource provisioning and scheduling in cloud environments. Please visit his homepage: <http://adelnadjarantoosi.info>.



Richard O. Sinnott is the Director of eResearch at the University of Melbourne and Professor of Applied Computing Systems. In these roles he is responsible for all aspects of eResearch (research-oriented IT development) at the University. He has been lead software engineer/architect on an extensive portfolio of national and international projects, with specific focus on those research domains requiring finer-grained access control (security). Prior to coming to Melbourne, Richard was the Technical Director of the UK National e-Science Centre; Director of e-Science at the University of Glasgow; Deputy Director (Technical) for the Bioinformatics Research Centre also at the University of Glasgow, and for a while the Technical Director of the National Centre for e-Social Science. He has a Ph.D. in Computing Science, an M.Sc. in Software Engineering and a B.Sc. in Theoretical Physics (Hons). He has over 300 peer-reviewed publications across a range of computing and application-specific domains. He teaches High Performance Computing and Cloud Computing at the University of Melbourne.



Rajkumar Buyya is a Fellow of IEEE, Professor of Computer Science and Software Engineering and Director of the Cloud Computing and Distributed Systems (CLOUDS) Laboratory at the University of Melbourne, Australia. He is also serving as the founding CEO of Manjrasoft, a spin-off company of the University, commercializing its innovations in Cloud Computing. He served as a Future Fellow of the Australian Research Council during 2012–2016. He has authored over 525 publications and seven text books including “Mastering Cloud Computing” published by McGraw Hill, China Machine Press, and Morgan Kaufmann for Indian, Chinese and international markets respectively. He also edited several books including “Cloud Computing: Principles and Paradigms” (Wiley Press, USA, Feb 2011). He is one of the highly cited authors in computer science and software engineering worldwide (h-index = 108, g-index = 225, 55, 200+ citations). Recently, Dr. Buyya is recognized as “2016 Web of Science Highly Cited Researcher” by Thomson Reuters. Software technologies for Grid and Cloud computing developed under Dr. Buyya's leadership have gained rapid acceptance and are in use at several academic institutions and commercial enterprises in 40 countries around the world. Manjrasoft's Aneka Cloud technology developed under his leadership has received “2010 Frost & Sullivan New Product Innovation Award”. Recently, Dr. Buyya received “Bharath Nirman Award” and “Mahatma Gandhi Award” along with Gold Medals for his outstanding and extraordinary achievements in Information Technology field and services rendered to promote greater friendship and India–International cooperation. He served as the founding Editor-in-Chief of the IEEE Transactions on Cloud Computing. He is currently serving as Co-Editor-in-Chief of Journal of Software: Practice and Experience, which was established over 45 years ago. For further information on Dr. Buyya, please visit his cyberhome: www.buyya.com.