



Big Data computing and clouds: Trends and future directions



Marcos D. Assunção^{a,*}, Rodrigo N. Calheiros^b, Silvia Bianchi^c, Marco A.S. Netto^c,
Rajkumar Buyya^{b,*}

^a INRIA, LIP, ENS de Lyon, France

^b The University of Melbourne, Australia

^c IBM Research, Brazil

HIGHLIGHTS

- Survey of solutions for carrying out analytics and Big Data on Clouds.
- Identification of gaps in technology for Cloud-based analytics.
- Recommendations of research directions for Cloud-based analytics and Big Data.

ARTICLE INFO

Article history:

Received 25 October 2013

Received in revised form

20 May 2014

Accepted 18 August 2014

Available online 27 August 2014

Keywords:

Big Data

Cloud computing

Analytics

Data management

ABSTRACT

This paper discusses approaches and environments for carrying out analytics on Clouds for Big Data applications. It revolves around four important areas of analytics and Big Data, namely (i) data management and supporting architectures; (ii) model development and scoring; (iii) visualisation and user interaction; and (iv) business models. Through a detailed survey, we identify possible gaps in technology and provide recommendations for the research community on future directions on Cloud-supported Big Data computing and analytics solutions.

© 2014 Elsevier Inc. All rights reserved.

1. Introduction

Society is becoming increasingly more instrumented and as a result, organisations are producing and storing vast amounts of data. Managing and gaining insights from the produced data is a challenge and key to competitive advantage. Analytics solutions that mine structured and unstructured data are important as they can help organisations gain insights not only from their privately acquired data, but also from large amounts of data publicly available on the Web [118]. The ability to cross-relate private information on consumer preferences and products with information from tweets, blogs, product evaluations, and data from social networks opens a wide range of possibilities for organisations to understand the needs of their customers, predict their wants and demands, and optimise the use of resources. This paradigm is being popularly termed as Big Data.

Despite the popularity on analytics and Big Data, putting them into practice is still a complex and time consuming endeavour. As Yu [136] points out, Big Data offers substantial value to organisations willing to adopt it, but at the same time poses a considerable number of challenges for the realisation of such added value. An organisation willing to use analytics technology frequently acquires expensive software licences; employs large computing infrastructure; and pays for consulting hours of analysts who work with the organisation to better understand its business, organise its data, and integrate it for analytics [120]. This joint effort of organisation and analysts often aims to help the organisation understand its customers' needs, behaviours, and future demands for new products or marketing strategies. Such effort, however, is generally costly and often lacks flexibility. Nevertheless, research and application of Big Data are being extensively explored by governments, as evidenced by initiatives from USA [20] and UK [106]; by academics, such as the bigdata@csail initiative from MIT [19]; and by companies such as Intel [122].

Cloud computing has been revolutionising the IT industry by adding flexibility to the way IT is consumed, enabling organisations to pay only for the resources and services they use. In an effort to

* Corresponding authors.

E-mail addresses: assuncao@acm.org (M.D. Assunção), rbuyya@unimelb.edu.au (R. Buyya).

reduce IT capital and operational expenditures, organisations of all sizes are using Clouds to provide the resources required to run their applications. Clouds vary significantly in their specific technologies and implementation, but often provide infrastructure, platform, and software resources as services [25,13].

The most often claimed benefits of Clouds include offering resources in a pay-as-you-go fashion, improved availability and elasticity, and cost reduction. Clouds can prevent organisations from spending money for maintaining peak-provisioned IT infrastructure that they are unlikely to use most of the time. Whilst at first glance the value proposition of Clouds as a platform to carry out analytics is strong, there are many challenges that need to be overcome to make Clouds an ideal platform for scalable analytics.

In this article we survey approaches, environments, and technologies on areas that are key to Big Data analytics capabilities and discuss how they help building analytics solutions for Clouds. We focus on the most important technical issues on enabling Cloud analytics, but also highlight some of the non-technical challenges faced by organisations that want to provide analytics as a service in the Cloud. In addition, we describe a set of gaps and recommendations for the research community on future directions on Cloud-supported Big Data computing.

2. Background and methodology

Organisations are increasingly generating large volumes of data as result of instrumented business processes, monitoring of user activity [14,127], web site tracking, sensors, finance, accounting, among other reasons. With the advent of social network Web sites, users create records of their lives by daily posting details of activities they perform, events they attend, places they visit, pictures they take, and things they enjoy and want. This data deluge is often referred to as Big Data [99,55,17]; a term that conveys the challenges it poses on existing infrastructure with respect to storage, management, interoperability, governance, and analysis of the data.

In today's competitive market, being able to explore data to understand customer behaviour, segment customer base, offer customised services, and gain insights from data provided by multiple sources is key to competitive advantage. Although decision makers would like to base their decisions and actions on insights gained from this data [43], making sense of data, extracting non obvious patterns, and using these patterns to predict future behaviour are not new topics. Knowledge Discovery in Data (KDD) [50] aims to extract non obvious information using careful and detailed analysis and interpretation. Data mining [133,84], more specifically, aims to discover previously unknown interrelations among apparently unrelated attributes of data sets by applying methods from several areas including machine learning, database systems, and statistics. Analytics comprises techniques of KDD, data mining, text mining, statistical and quantitative analysis, explanatory and predictive models, and advanced and interactive visualisation to drive decisions and actions [43,42,63].

Fig. 1 depicts the common phases of a traditional analytics workflow for Big Data. Data from various sources, including databases, streams, marts, and data warehouses, are used to build models. The large volume and different types of the data can demand pre-processing tasks for integrating the data, cleaning it, and filtering it. The prepared data is used to train a model and to estimate its parameters. Once the model is estimated, it should be validated before its consumption. Normally this phase requires the use of the original input data and specific methods to validate the created model. Finally, the model is consumed and applied to data as it arrives. This phase, called model scoring, is used to generate predictions, prescriptions, and recommendations. The results are interpreted and evaluated, used to generate new models or calibrate existing ones, or are integrated to pre-processed data.

Analytics solutions can be classified as descriptive, predictive, or prescriptive as illustrated in Fig. 2. Descriptive analytics uses historical data to identify patterns and create management reports; it is concerned with modelling past behaviour. Predictive analytics attempts to predict the future by analysing current and historical data. Prescriptive solutions assist analysts in decisions by determining actions and assessing their impact regarding business objectives, requirements, and constraints.

Despite the hype about it, using analytics is still a labour intensive endeavour. This is because current solutions for analytics are often based on proprietary appliances or software systems built for general purposes. Thus, significant effort is needed to tailor such solutions to the specific needs of the organisation, which includes integrating different data sources and deploying the software on the company's hardware (or, in the case of appliances, integrating the appliance hardware with the rest of the company's systems) [120]. Such solutions are usually developed and hosted on the customer's premises, are generally complex, and their operations can take hours to execute. Cloud computing provides an interesting model for analytics, where solutions can be hosted on the Cloud and consumed by customers in a pay-as-you-go fashion. For this delivery model to become reality, however, several technical issues must be addressed, such as data management, tuning of models, privacy, data quality, and data currency.

This work highlights technical issues and surveys existing work on solutions to provide analytics capabilities for Big Data on the Cloud. Considering the traditional analytics workflow presented in Fig. 1, we focus on key issues in the phases of an analytics solution. With Big Data it is evident that many of the challenges of Cloud analytics concern data management, integration, and processing. Previous work has focused on issues such as data formats, data representation, storage, access, privacy, and data quality. Section 3 presents existing work addressing these challenges on Cloud environments. In Section 4, we elaborate on existing models to provide and evaluate data models on the Cloud. Section 5 describes solutions for data visualisation and customer interaction with analytics solutions provided by a Cloud. We also highlight some of the business challenges posed by this delivery model when we discuss service structures, service level agreements, and business models. Security is certainly a key challenge for hosting analytics solutions on public Clouds. We consider, however, that security is an extensive topic and would hence deserve a study of its own. Therefore, security and evaluation of data correctness [130] are out of scope of this survey.

3. Data management

One of the most time-consuming and labour-intensive tasks of analytics is preparation of data for analysis; a problem often exacerbated by Big Data as it stretches existing infrastructure to its limits. Performing analytics on large volumes of data requires efficient methods to store, filter, transform, and retrieve the data. Some of the challenges of deploying data management solutions on Cloud environments have been known for some time [1,113,82], and solutions to perform analytics on the Cloud face similar challenges. Cloud analytics solutions need to consider the multiple Cloud deployment models adopted by enterprises, where Clouds can be for instance:

- *Private*: deployed on a private network, managed by the organisation itself or by a third party. A private Cloud is suitable for businesses that require the highest level of control of security and data privacy. In such conditions, this type of Cloud infrastructure can be used to share the services and data more efficiently across the different departments of a large enterprise.
- *Public*: deployed off-site over the Internet and available to the general public. Public Cloud offers high efficiency and shared

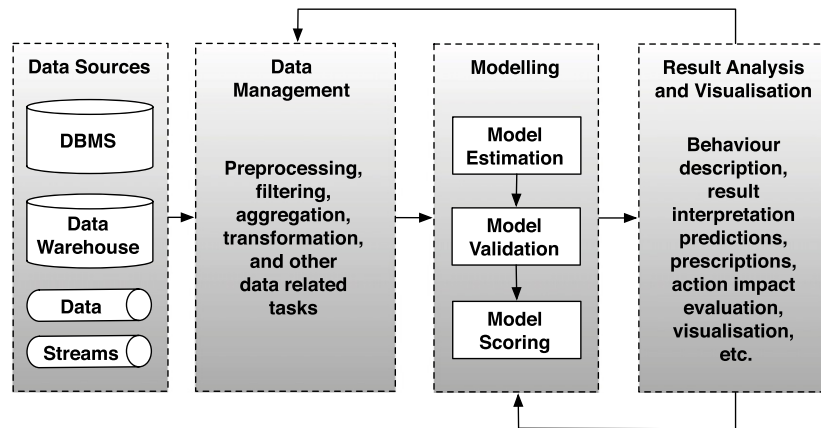


Fig. 1. Overview of the analytics workflow for Big Data.




Fig. 2. Categories of analytics.

resources with low cost. The analytics services and data management are handled by the provider and the quality of service (e.g. privacy, security, and availability) is specified in a contract. Organisations can leverage these Clouds to carry out analytics with a reduced cost or share insights of public analytics results.

- **Hybrid:** combines both Clouds where additional resources from a public Cloud can be provided as needed to a private Cloud. Customers can develop and deploy analytics applications using a private environment, thus reaping benefits from elasticity and higher degree of security than using only a public Cloud.

Considering the Cloud deployments, the following scenarios are generally envisioned regarding the availability of data and analytics models [87]: (i) data and models are private; (ii) data is public, models are private; (iii) data and models are public; and (iv) data is private, models are public. Jensen et al. [79] advocate on deployment models for Cloud analytics solutions that vary from solutions using privately hosted software and infrastructure, to private analytics hosted on a third party infrastructure, to public model where the solutions are hosted on a public Cloud.

Different from traditional Cloud services, analytics deals with high-level capabilities that often demand very specialised resources such as data and domain experts' analysis skills. For this reason, we advocate that under certain business models – especially those where data and models reside on the provider's premises – not only ordinary Cloud services, but also the skills of data experts need to be managed. To achieve economies of scale and elasticity, Cloud-enabled Big Data analytics needs to explore means to allocate and utilise these specialised resources in a proper manner. The rest of this section discusses existing solutions on data management irrespective of where data experts are physically located, focusing on storage and retrieval of data for analytics; data diversity, velocity and integration; and resource scheduling for data processing tasks.

3.1. Data variety and velocity

Big Data is characterised by what is often referred to as a multi-V model, as depicted in Fig. 3. Variety represents the data types, velocity refers to the rate at which the data is produced and processed, and volume defines the amount of data. Veracity refers to how much the data can be trusted given the reliability of its source [136], whereas value corresponds the monetary worth that




Fig. 3. Some 'Vs' of Big Data.

a company can derive from employing Big Data computing. Although the choice of Vs used to explain Big Data is often arbitrary and varies across reports and articles on the Web – e.g. as of writing Viability is becoming a new V – variety, velocity, and volume [112,140] are the items most commonly mentioned.

Regarding Variety, it can be observed that over the years, substantial amount of data has been made publicly available for scientific and business uses. Examples include repositories with government statistics¹; historical weather information and forecasts; DNA sequencing; information on traffic conditions in large metropolitan areas; product reviews and comments; demographics [105]; comments, pictures, and videos posted on social network Web sites; information gathered using citizen-science platforms [22]; and data collected by a multitude of sensors measuring various environmental conditions such as temperature, air humidity, air quality, and precipitation.

An example illustrating the need for such a variety within a single analytics application is the Eco-Intelligence [139] platform. Eco-Intelligence was designed to analyse large amounts of data to support city planning and promote more sustainable development. The platform aims to efficiently discover and process data from several sources, including sensors, news, Web sites, television and radio, and exploit information to help urban stakeholders cope with the highly dynamics of urban development. In a related scenario, the Mobile Data Challenge (MDC) was created aimed at generating innovations on smartphone-based research, and to enable community evaluation of mobile data analysis methodologies [90]. Data from around 200 users of mobile phones was collected over a year as part of the Lausanne Data Collection Campaign. Another related area benefiting from analytics is Massively Multiplayer Online Game (MMOGs). CAMEO [78] is an architecture for continuous analytics for MMOGs that uses Cloud resources for analysis of tasks. The architecture provides mechanisms for data collection and continuous analytics on several factors such as understanding the needs of the game community.

¹ <http://www.data.gov>.




Fig. 4. Variety of data.




Fig. 5. Velocity of data.

Data is also often available for sale in the format of research and technical reports, market segment and financial analyses, among other means. This data can be used by various applications, for instance, to improve the living conditions in large cities, to provide better quality services, to optimise the use of natural resources,² and to prevent or manage response to unplanned events.

Handling and analysing this data poses several challenges as it can be of different types (Fig. 4). It is argued that a large part of data produced today is either *unstructured* or *semi-structured*.

Considering data velocity, it is noticed that, to complicate matters further, data can arrive and require processing at different speeds, as illustrated in Fig. 5. Whilst for some applications, the arrival and processing of data can be performed in batch, other analytics applications require continuous and real-time analyses, sometimes requiring immediate action upon processing of incoming data streams. For instance, to provide active management for data centres, Wang et al. [131] present an architecture that integrates monitoring and analytics. The proposed architecture relies on Distributed Computation Graphs (DCG) that are created to implement the desired analytics functions. The motivating use cases consist in scenarios where information can be collected from monitored equipments and services, and once a potential problem is identified, the system can instantiate DCGs to collect further information for analytics.

Increasingly often, data arriving via streams needs to be analysed and compared against historical information. Different data sources may use their own formats, which makes it difficult to *integrate* data from multiple sources in an analytics solution. As highlighted in existing work [52], *standard formats and interfaces* are crucial so that solution providers can benefit from economies of scale derived from data integration capabilities that address the needs of a wide range of customers.

The rest of our discussion on data management for Cloud analytics surrounds these two Vs of Big Data, namely Variety and Velocity. We survey solutions on how this diverse data is stored, how it can be integrated and how it is often processed. The discussion on visualisation also explores Velocity by highlighting factors such as interactivity and batch based visualisation. Although the other Vs of Big Data are important, we consider that some of them, as discussed earlier, deserve a study of their own, such as data Veracity. Other Vs are subjective; Volume is highly dependent on the scalability of existing hardware infrastructure, which improves quickly and can render a survey obsolete very rapidly; Value may depend on how efficient a company employs the analytics solutions at hand. Besides the V attributes, Big Data analytics also shares concerns with other data-related disciplines, and thus can directly benefit from the body of knowledge developed in the last years on such established subjects. This is the case of issues such as data quality [110] and data provenance [102].

3.2. Data storage

Several solutions were proposed to store and retrieve large amounts of data demanded by Big Data, some of which are currently used in Clouds. Internet-scale file systems such as the Google File System (GFS) [57] attempt to provide the robustness, scalability, and reliability that certain Internet services need. Other solutions provide object-store capabilities where files can be replicated across multiple geographical sites to improve redundancy, scalability, and data availability. Examples include Amazon Simple Storage Service (S3),³ Nirvanix Cloud Storage,⁴ OpenStack Swift⁵ and Windows Azure Binary Large Object (Blob) storage.⁶ Although these solutions provide the scalability and redundancy that many Cloud applications require, they sometimes do not meet the concurrency and performance needs of certain analytics applications.

One key aspect in providing performance for Big Data analytics applications is the *data locality*. This is because the volume of data involved in the analytics makes it prohibitive to transfer the data to process it. This was the preferred option in typical high performance computing systems: in such systems, that typically concern performing CPU-intensive calculations over a moderate to medium volume of data, it is feasible to transfer data to the computing units, because the ratio of data transfer to processing time is small. Nevertheless, in the context of Big Data, this approach of moving data to computation nodes would generate large ratio of data transfer time to processing time. Thus, a different approach is preferred, where computation is moved to where the data is. The same approach of exploring data locality was explored previously in scientific workflows [47] and in Data Grids [128].

In the context of Big Data analytics, MapReduce presents an interesting model where data locality is explored to improve the performance of applications. Hadoop, an open source MapReduce implementation, allows for the creation of clusters that use the Hadoop Distributed File System (HDFS) to partition and replicate data sets to nodes where they are more likely to be consumed by mappers. In addition to exploiting concurrency of large numbers of nodes, HDFS minimises the impact of failures by replicating data sets to a configurable number of nodes. It has been used by Thussu et al. [125] to develop an analytics platform to process Facebook's large data sets. The platform uses Scribe to aggregate logs from Web servers and then exports them to HDFS files and uses a Hive-Hadoop cluster to execute analytics jobs. The platform includes replication and compression techniques and columnar compression of Hive⁷ to store large amounts of data.

Among the drawbacks of Cloud storage techniques and MapReduce implementations, there is the fact that they require the customer to learn a new set of APIs to build analytics solutions for the Cloud. To minimise this hurdle, previous work has also investigated POSIX-like file systems for data analytics. As an example, Ananthanarayanan et al. [6] adapted POSIX-based cluster file systems to be used as data storage for Cloud analytics applications. By using the concept of meta-blocks, they demonstrated that IBM's General Parallel File System (GPFS) [117] can match the read performance of HDFS. A meta-block is a consecutive set of data blocks that are allocated in the same disk, thus guaranteeing contiguity. The proposed approach explores the trade-off between different block sizes, where meta-blocks minimise seek overhead in MapReduce applications, whereas small blocks reduce pre-fetch overhead and improves cache management for ordinary applications.

³ <http://aws.amazon.com/s3/>.

⁴ <http://www.nirvanix.com>.

⁵ <http://swift.openstack.org>.

⁶ <http://www.windowsazure.com/en-US/services/data-management/>.

⁷ <http://hive.apache.org>.

² Sense-T laboratory: <http://www.sense-t.org.au/about/the-big-picture>.

Tantisiroj et al. [124] compared the Parallel Virtual File System (PVFS) [28] against HDFS, where they observed that PVFS did not present significant improvement in completion time and throughput compared to HDFS.

Although a large part of the data produced nowadays is unstructured, relational databases have been the choice most organisations have made to store data about their customers, sales, and products, among other things. As data managed by traditional DBMS ages, it is moved to data warehouses for analysis and for sporadic retrieval. Models such as MapReduce are generally not the most appropriate to analyse such relational data. Attempts have been made to provide hybrid solutions that incorporate MapReduce to perform some of the queries and data processing required by DBMS's [1]. Cohen et al. [39] provide a parallel database design for analytics that supports SQL and MapReduce scripting on top of a DBMS to integrate multiple data sources. A few providers of analytics and data mining solutions, by exploring models such as MapReduce, are migrating some of the processing tasks closer to where the data is stored, thus trying to minimise surpluses of siloed data preparation, storage, and processing [85]. Data processing and analytics capabilities are moving towards Enterprise Data Warehouses (EDWs), or are being deployed in data hubs [79] to facilitate reuse across various data sets.

With respect to EDW, some Cloud providers offer solutions that promise to scale to one petabyte of data or more. Amazon Redshift [2], for instance, offers columnar storage and data compression and aims to deliver high query performance by exploring a series of features, including a massively parallel processing architecture using high performance hardware, mesh networks, locally attached storage, and zone maps to reduce the I/O required by queries. Amazon Data Pipeline [3] allows a customer to move data across different Amazon Web Services, such as Elastic MapReduce (EMR) [4] and DynamoDB [46], and hence compose the required analytics capabilities.

Another distinctive trend in Cloud computing is the increasing use of NoSQL databases as the preferred method for storing and retrieving information. NoSQL adopts a non-relational model for data storage. Leavitt argues that non-relational models have been available for more than 50 years in forms such as object-oriented, hierarchical, and graph databases, but recently this paradigm started to attract more attention with models such as key-store, column-oriented, and document-based stores [92]. The causes for such raise in interest, according to Levitt, are better performance, capacity of handling unstructured data, and suitability for distributed environments [92].

Han et al. [68] presented a survey of NoSQL databases with emphasis on their advantages and limitations for Cloud computing. The survey classifies NoSQL systems according to their capacity in addressing different pairs of CAP (consistency, availability, partitioning). The survey also explores the data model that the studied NoSQL systems support.

Hecht and Jablonski [69] compared different NoSQL systems in regard to supported data models, types of query supported, and support for concurrency, consistency, replication, and partitioning. Hecht and Jablonski concluded that there are big differences among the features of different technologies, and there is no single system that would be the most suitable for every need. Therefore, it is important for adopters to understand the requirements of their applications and the capabilities of different systems so that the system whose features better match their needs is selected [69].

3.3. Data integration solutions

Forrester Research published a technical report that discusses some of the problems that traditional Business Intelligence (BI) faces [85], highlighting that there is often a surplus of siloed data

preparation, storage, and processing. Authors of the report envision some data processing and Big Data analytics capabilities being migrated to the EDW, hence freeing organisations from unnecessary data transfer and replication and the use of disparate data-processing and analysis solutions. Moreover, as discussed earlier, they advocated that analytics solutions will increasingly expose data processing and analysis features via MapReduce and SQL–MR-like interfaces. SAP HANA One [115], as an example, is an in-memory platform hosted by Amazon Web Services that provides real-time analytics for SAP applications. HANA One also offers a SAP data integrator to load data from HDFS and Hive-accessible databases.

EDWs or Cloud based data warehouses, however, create certain issues with respect to data integration and the addition of new data sources. Standard formats and interfaces can be essential to achieve economies of scale and meet the needs of a large number of customers [52]. Some solutions attempt to address some of these issues [105,21]. Birst [21] provides composite spaces and space inheritance, where a composite space integrates data from one or more parent spaces with additional data added to the composite space. Birst provides a Software as a Service (SaaS) solution that offers analytics functionalities on a subscription model; and appliances with the business analytics infrastructure, hence providing a model that allows a customer to migrate gradually from an on-premise analytics to a scenario with Cloud-provided analytics infrastructure. To improve the market penetration of analytics solutions in emerging markets such as India, Deepak et al. [48] propose a multi-flow solution for analytics that can be deployed on the Cloud. The multi-flow approach provides a range of possible analytics operators and flows to compose analytics solutions; viewed as workflows or instantiations of a multi-flow solution. IVOCA [18] is a tool aimed at Customer Relationship Management (CRM) that ingests both structured and unstructured data and provides data linking, classification, and text mining tools to facilitate analysts' tasks and reduce the time to insight.

Habich et al. [67] propose Web services that co-ordinate data Clouds for exchanging massive data sets. The Business Process Execution Language (BPEL) data transition approach is used for data exchange by passing references to data between services to reduce the execution time and guarantee the correct data processing of an analytics process. A generic data Cloud layer is introduced to handle heterogeneous data Clouds, and is responsible for mapping generic operations to each Cloud implementation. DataDirect Cloud [41] also provides generic interfaces by offering JDBC/ODBC drivers for applications to execute SQL queries against different databases stored on a Cloud. Users are not required to deal with different APIs and query languages specific to each Cloud storage solution.

PivotLink's AnalyticsCLOUD [105] handles both structured and unstructured data, providing data integration features. PivotLink also provides DataCLOUD with information about over 350 demographic, hobbies, and interest data fields for 120 million US households. This information can be used by customers to perform brand sentiment analysis [51] and verify how weather affects their product performance.

3.4. Data processing and resource management

MapReduce [45] is one of the most popular programming models to process large amounts of data on clusters of computers. Hadoop [10] is the most used open source MapReduce implementation, also made available by several Cloud providers [4,16,77,132]. Amazon EMR [4] enables customers to instantiate Hadoop clusters to process large amounts of data using the Amazon Elastic Compute Cloud (EC2) and other Amazon Web Services for data storage and transfer.

Hadoop uses the HDFS file system to partition and replicate data sets across multiple nodes, such that when running a MapReduce application, a mapper is likely to access data that is locally stored on the cluster node where it is executing. Although Hadoop provides a set of APIs that allows developers to implement MapReduce applications, very often a Hadoop workflow is composed of jobs that use high-level query languages such as Hive and Pig Latin, created to facilitate search and specification of processing tasks. Lee et al. [94] present a survey about the features, benefits, and limitations of MapReduce for parallel data analytics. They also discuss extensions proposed for this programming model to overcome some of its limitations.

Hadoop provides data parallelism and its data and task replication schemes enable fault tolerance, but what is often criticised about it is the time required to load data into HDFS and the lack of reuse of data produced by mappers. MapReduce is a model created to exploit commodity hardware, but when executed on reliable infrastructure, the mechanisms it provides to deal with failures may not be entirely essential. Some of the provided features can be disabled in certain scenarios. Herodotou and Babu [71] present techniques for profiling MapReduce applications, identifying bottlenecks and simulating what-if scenarios. Previous work has also proposed optimisations to handle these shortcomings [66]. Cuzzocrea et al. [40] discuss issues concerning analytics over big multidimensional data and the difficulties in building multidimensional structures in HDFS and integrating multiple data sources to Hadoop.

Starfish [72], a data analytics system built atop Hadoop, focuses on improving the performance of clusters throughout the data lifecycle in analytics, without requiring users to understand the available configuration options. Starfish employs techniques at several levels to optimise the execution of MapReduce jobs. It uses dynamic instrumentation to profile jobs and optimises workflows by minimising the impact of data unbalance and by balancing the load of executions. Starfish's Elastisizer automates provisioning decisions using a mix of simulation and model-based estimation to address what-if questions on workload performance.

Lee et al. [93] present an approach that allocates resources and schedules jobs considering data analytics workloads, in order to enable consolidation of a cluster workload, reducing the number of machines allocated for processing the workload during periods of small load. The approach uses Hadoop and works with two pools of machines – *core* and *accelerator* – and dynamically adjusts the size of each pool according to the observed load.

Daytona [16], a MapReduce runtime for Windows Azure, leverages the scalable storage services provided by Azure's Cloud infrastructure as the source and destination of data. It uses Cloud features to provide load balancing and fault tolerance. The system relies on a master-slave architecture where the master is responsible for scheduling tasks and the slaves for carrying out map and reduce operations. Section 5 discusses the visualisation features that Daytona provides.

Previous work shows that there is an emerging class of MapReduce applications that feature small, short, and highly interactive jobs [31,54]. As highlighted in Section 5, the visualisation community often criticises the lack of interactivity of MapReduce-based analytics solutions. Over the past few years, however, several attempts have been made to tackle this issue. Borthakur et al. [23], for instance, describe optimisations implemented in HDFS and HBase⁸ to make them more responsive to the realtime requirements of Facebook applications. Chen et al. [30] propose energy efficiency improvements to Hadoop by maintaining two distinct pools of resources, namely interactive and batch jobs.

The eXtreme Analytics Platform (XAP) [15] enables analytics supporting multiple data sources, data types (structured and unstructured), and multiple types of analyses. The target infrastructure of the architecture is a cluster running a distributed file system. A modified version of Hadoop, deployed in the cluster, contains an application scheduler (FLEX) able to better utilise the available resources than the default Hadoop scheduler. The analytics jobs are created via a high-level language script, called Jaql, that converts the high-level descriptive input into an analytics MapReduce workflow that is executed in the target infrastructure.

Previous work has also considered other models for performing analytics, such as scientific workflows and Online Analytical Processing (OLAP). Rahman et al. [108] propose a hybrid heuristic for scheduling data analytics workflows on heterogeneous Cloud environments; a heuristic that optimises cost of workflow execution and satisfies users requirements, such as budget, deadline, and data placement.

In the field of simulation-enabled analytics, Li et al. [97] developed an analytical application, modelled as a Direct Acyclic Graph (DAG), for predicting the spread of dengue fever outbreaks in Singapore. The analytics workflow receives data from multiple sources, including current and past data about climate and weather from meteorological agencies and historical information about dengue outbreaks in the country. This data, with user-supplied input about the origin of the infection, is used to generate a map of the spread of the disease in the country in a day-by-day basis. A hybrid Cloud is used to speed up the application execution. Other characteristics of the application are security features and cost-effective exploration of Cloud resources: the system keeps the utilisation of public Cloud resources to a minimum to enable the analytics to complete in the specified time and budget. A public Cloud has also been used in a similar scenario to simulate the impact of public transport disruptions on urban mobility [81].

Chohan et al. [36] evaluated the support of OLAP for Google App Engine (GAE) [58] highlighting limitations and assessing their impact on cost and performance of applications. A hybrid approach to perform OLAP using GAE and AppScale [24] was provided, using two methods for data synchronisation, namely *bulk data transfer* and *incremental data transfer*. Moreover, Jung et al. [80] propose optimisations for scheduling and processing of Big Data analysis on federated Clouds.

Chang et al. [29] examined different data analytics workloads, where results show significant diversity of resource usage (CPU, I/O and, network). They recommend the use of transformation mechanisms such as indexing, compression, and approximation to provide a balanced system and improve efficiency of data analysis.

The Cloud can also be used to extend the capabilities of analyses initially started on the customer's premises. CloudComet, for example, is an autonomic computing engine that supports Cloud bursts that has been used to provide the programming and runtime infrastructure to scale out/in certain on-line risk analyses [83]. CloudComet and commercial technologies such as Aneka [27] can utilise both private resources and resources from a public Cloud provider to handle peaks in the demands of online risk analytics.

Some analytics applications including stock quotes and weather prediction have stringent time constraints, usually falling in the near-time and stream categories described earlier. Request processing time is important to deliver results in a timely fashion. Chen et al. [33] investigate Continuous Analytics as a Service (CaaS) that blends stream processing and relational data techniques to extend the DBMS model and enable real-time continuous analytics service provisioning. The dynamic stream processing and static data management for data intensive analytics are unified by providing an SQL-like interface to access both static and stream data. The proposed cycle-based query model and transaction model allow SQL queries to run and to commit per cycle whilst analysing stream data per chunk. The analysis results are made visible to clients whilst a continued query for results generation

⁸ <http://hbase.apache.org>.

Table 1
Summary of works on model building and scoring.

Work	Goal	Service model	Deployment model
Guazzelli et al. [64]	Predictive analytics (scoring)	IaaS	Public
Zementis [138]	Data analysis and model building	SaaS	Public or private
Google Prediction API [59]	Model building	SaaS	Public
Apache Mahout [11]	Data analysis and model building	IaaS	Any
Hazy [88]	Model building	IaaS	Any

is still running. Existing work on stream and near-time processing attempt to leverage strategies to predict user or service behaviour [137]. In this way, an analytics service can pre-fetch data to anticipate a user's behaviour, hence selecting the appropriate applications and methods before the user's request arrives.

Realtime analysis of Big Data is a hot topic, with Cloud providers increasingly offering solutions that can be used as building blocks of stream and complex event processing systems. AWS Kinesis [5] is an elastic system for real-time processing of streaming data that can handle multiple sources, be used to build dashboards, handle events, and generate alerts. It allows for integration with other AWS services. In addition, stream processing frameworks including Apache S4 [9], Storm [121] and IBM InfoSphere Streams [75] can be deployed on existing Cloud offerings. Software systems such as *storm-deploy*, a Clojure project based on Pallet,⁹ aim to ease deployment of Storm topologies on Cloud offerings including AWS EC2. Suro, a data pipeline system used by Netflix to collect events generated by its applications, has recently been made available to the broader community as an open source project [8]. Aiming to address similar requirements, Apache Kafka [62] is a real-time publish-subscribe infrastructure initially used at LinkedIn to process activity data and later released as an open source project. Incubated by the Apache Software Foundation, Samza [12] is a distributed stream processing framework that blends Kafka and Apache Hadoop YARN. Whilst Samza provides a model where streams are the input and output to jobs, execution is completely handled by YARN.

3.5. Challenges in big data management

In this section, we discuss current research targeting the issue of Big Data management for analytics. There are still, however, many open challenges in this topic. The list below is not exhaustive, and as more research in this field is conducted, more challenging issues will arise.

Data variety: How to handle an always increasing volume of data? Especially when the data is unstructured, how to quickly extract meaningful content out of it? How to aggregate and correlate streaming data from multiple sources?

Data storage: How to efficiently recognise and store important information extracted from unstructured data? How to store large volumes of information in a way it can be timely retrieved? Are current file systems optimised for the volume and variety demanded by analytics applications? If not, what new capabilities are needed? How to store information in a way that it can be easily migrated/ported between data centres/Cloud providers?

Data integration: New protocols and interfaces for integration of data that are able to manage data of different nature (structured, unstructured, semi-structured) and sources.

Data Processing and Resource Management: New programming models optimised for streaming and/or multidimensional data; new backend engines that manage optimised

file systems; engines able to combine applications from multiple programming models (e.g. MapReduce, workflows, and bag-of-tasks) on a single solution/abstraction. How to optimise resource usage and energy consumption when executing the analytics application?

4. Model building and scoring

The data storage and Data as a Service (DaaS) capabilities provided by Clouds are important, but for analytics, it is equally relevant to use the data to build models that can be utilised for forecasts and prescriptions. Moreover, as models are built based on the available data, they need to be tested against new data in order to evaluate their ability to forecast future behaviour. Existing work has discussed means to offload such activities – termed here as model building and scoring – to Cloud providers and ways to parallelise certain machine learning algorithms [126,11,74]. This section describes work on the topic. Table 1 summarises the analysed work, its goals, and target infrastructures.

Guazzelli et al. [64] use Amazon EC2 as a hosting platform for the Zementis' ADAPA model [138] scoring engine. Predictive models, expressed in Predictive Model Markup Language (PMML) [65], are deployed in the Cloud and exposed via Web Services interfaces. Users can access the models with Web browser technologies to compose their data mining solutions. Existing work also advocates the use of PMML as a language to exchange information about predictive models [73].

Zementis [138] also provides technologies for data analysis and model building that can run either on a customer's premises or be allocated as SaaS using Infrastructure as a Service (IaaS) provided by solutions such as Amazon EC2 and IBM SmartCloud Enterprise [76].

Google Prediction API [59] allows users to create machine learning models to predict numeric values for a new item based on values of previously submitted training data or predict a category that best describes an item. The prediction API allows users to submit training data as comma separated files following certain conventions, create models, share their models or use models that others shared. With the Google Prediction API, users can develop applications to perform analytics tasks such as sentiment analysis [51], purchase prediction, provide recommendations, analyse churn, and detect spam. The Apache Mahout project [11] aims to provide tools to build scalable machine learning libraries on top of Hadoop using the MapReduce paradigm. The provided libraries can be deployed on a Cloud and be explored to build solutions that require clustering, recommendation mining, document categorisation, among others.

By trying to ease the complexity in building trained systems such as IBM's Watson, Apple's Siri and Google Knowledge Graph, the Hazy project [88] focuses on identifying and validating two categories of abstractions in building trained systems, namely *programming abstractions* and *infrastructure abstractions*. It is argued that, by providing such abstractions, it would be easier for one to assemble existing solutions and build trained systems. To achieve a small and compoundable programming interface, Hazy employs a data model that combines the relational data model and a probabilistic rule-based language. For infrastructure abstraction, Hazy

⁹ <http://github.com/pallet/pallet>.

leverages the observation that many statistical analysis algorithms behave as a user-defined aggregate in a Relational Database Management System (RDBMS). Hazy then explores features of the underlying infrastructure to improve the performance on these aggregates.

4.1. Open challenges

The key challenge in the area of Model Building and Scoring is the discovery of techniques that are able to explore the rapid elasticity and large scale of Cloud systems. Given that the amount of data available for Big Data analytics is increasing, timely processing of such data for building and scoring would give a relevant advantage for businesses able to explore such a capability.

In the same direction, standards and interfaces for these activities are also required, as they would help to disseminate “prediction and analytics as services” providers that would compete for customers. If the use of such services does not incur vendor lock in (via utilisation of standards APIs and formats), customers can choose the service provider only based on cost and performance of services, enabling the emergence of a new competitive market.

5. Visualisation and user interaction

With the increasing amounts of data with which analyses need to cope, good visualisation tools are crucial. These tools should consider the quality of data and presentation to facilitate navigation [44]. The type of visualisation may need to be selected according to the amount of data to be displayed, to improve both displaying and performance. Visualisation can assist in the three major types of analytics: descriptive, predictive, and prescriptive. Many visualisation tools do not describe advanced aspects of analytics, but there has been an effort to explore visualisation to help on predictive and prescriptive analytics, using for instance sophisticated reports and storytelling [86]. A key aspect to be considered on visualisation and user interaction in the Cloud is that network is still a bottleneck in several scenarios [123]. Users ideally would like to visualise data processed in the Cloud having the same experience and feel as though data were processed locally. Some solutions have been tackling this requirement.

For example, as Fisher et al. [52] point out, many Cloud platforms available to process data analytics tasks still resemble the *batch-job* model used in the early times of the computing era. Users typically submit their jobs and wait until the execution is complete to download and analyse sample results to validate full runs. As this back and forth of data is not well supported by the Cloud, the authors issue a call to arms for both research and development of better interactive interfaces for Big Data analytics where users iteratively pose queries and see rapid responses. Fisher et al. introduce *sampleAction* [53] to explore whether interactive techniques acting over only incremental samples can be considered as sufficiently trustworthy by analysts to make closer to real time decisions about their queries. Interviews with three teams of analysts suggest that representations of incremental query results were robust enough so that analysts were prepared either to abandon a query, refine it, or formulate new queries. King [84] also highlights the importance of making the analytics process iterative, with multiple checkpoints for assessment and adjustment.

In this line, existing work aims to explore the batch-job model provided by solutions including MapReduce as a backend to features provided in interfaces with which users are more familiar. Trying to leverage the popularity of spreadsheets as a tool to manipulate data and perform analysis, Barga et al. proposed an Excel ribbon connected to Daytona [16], a Cloud service for data storage and analytics. Users manipulate data sets on Excel and plugins use Microsoft’s Azure infrastructure [26] to run MapReduce applications. In addition, as described earlier, several improvements

have been proposed to MapReduce frameworks to handle interactive applications [23,30,100]. However, most of these solutions are not yet made available for general use in the Cloud.

Several projects attempt to provide a range of visualisation methods from which users can select a set that suits their requirements. ManyEyes [129] from IBM allows users to upload their data, select a visualisation method – varying from basic to advanced – and publish their results. Users may also navigate through existing visualisations and discuss their findings and experience with peers. Selecting data sources automatically or semi-automatically is also an important feature to help users perform analytics. PanXpan [104] is an example of a tool that automatically identifies the fields in structured data sets based on user analytics module selection. FusionCharts [56] is another tool to allow users to visually select a subset of data from the plotted data points to be submitted back to the server for further processing. CloudVista [35,135] is a software to help on visual data selection for further analysis refinement.

Existing work also provides means for users to aggregate data from multiple sources and employ various visualisation models, including dashboards, widgets, line and bar charts, demographics, among other models [105,98,60,61,103]. Some of these features can be leveraged to perform several tasks, including create reports; track what sections of a site are performing well and what kind of content can create better user experience; how information sharing on a social network impacts the web site usage; track mobile usage [14,127]; and evaluate the impact of advertising campaigns.

Choo and Park [37] argue that the reason why Big Data visualisation is not real time is the computational complexity of the analytics operations. In this direction, authors discuss strategies to reduce computational complexity of data analytics operations by, for instance, decreasing precision of calculations.

Apart from software optimisation, dedicated hardware for visualisation is becoming key for Big Data analytics. For example, Reda et al. [109] discuss that, although existing tools are able to provide data belonging to a range of classes, their dimensionality and volume exceed the capacity of visualisation provided by standard displays. This requires the utilisation of large-scale visualisation environments, such as CyberCommons and CAVE2, which are composed of a large display wall with resolution three orders of magnitude higher than that achieved by commercial displays [109]. Remote visualisation systems, such as Nautilus from XSEDE (Extreme Science and Engineering Discovery Environment—the new NSF TeraGrid project replacement), are becoming more common to supply high demand for memory and graphical processors to assist very large data visualisation [134].

Besides visualisation of raw data, summarised content in form of reports are essential to perform predictive and prescriptive analytics. Several solutions have explored report generation and visualisation. For instance, SAP Crystal Solutions [116] provides BI functionalities via which customers can explore available data to build reports with interactive charts, what-if scenarios, and dashboards. The produced reports can be visualised on the Web, e-mail, Microsoft Office, or be embedded into enterprise applications. Another example on report visualisation is Cloud9 Analytics [38], which aims to automate reports and dashboards, based on data from CRM and other systems. It provides features for sales reports, sales analytics, and sales forecasts and pipeline management. By exploring history data and using the notion of risk, it offers customers clues on which projects they should invest their resources and what projects or products require immediate action. Other companies also offer solutions that provide sales forecasts, change analytics, and customised reports [111,21]. Salesforce [114] supports customisable dashboards through collaborative analytics. The platform allows authorised users to share their charts and information with other users. Another trend on visualisation to help on predictive and prescriptive analytics is storytelling [86], which aims at presenting data with a narrative visualisation.

There are also visualisation tools specific for a given domain. For instance, in the field of climate-modelling, Lee et al. [95] developed a tool for visualisation of simulated Madden–Julian Oscillation, which is an important meteorological event that influences raining patterns from South America to Southeast Asia. The tool enables tracking of the event and its visualisation using Google Earth. In the area of computing networks management, Liao et al. [96] evaluated five approaches for visualisation of anomalies in large scale computer networks. Each method has its own applications depending on the specific type of anomaly to be visualised and the scale of the managed system. There are also solutions that provide means to visualise demographic information. Andrienko et al. [7] proposed interactive visual display for analysis of movement behaviour of people, vehicle, and animals. The visualisation tool displays the movement data, information about the time spent in a place, and the time interval from one place to another.

5.1. Open challenges

There are many research challenges in the field of Big Data visualisation. First, more efficient data processing techniques are required in order to enable real-time visualisation. Choo and Park [37] appoint some techniques that can be employed with this objective, such as reduction of accuracy of results, coarsely processing of data points, compatible with the resolution of the visualisation device, reduced convergence, and data scale confinement. Methods considering each of these techniques could be further researched and improved.

Cost-effective devices for large-scale visualisation is another hot topic for analytics visualisation, as they enable finer resolution than simple screens. Visualisation for management of computer networks and software analytics [101] is also an area that is attracting attention of researchers and practitioners for its extreme relevance to management of large-scale infrastructure (such as Clouds) and software, with implications in global software development, open source software development, and software quality improvements.

6. Business models and non-technical challenges

In addition to providing tools that customers can use to build their Big Data analytics solutions on the Cloud, models for delivering analytics capabilities as services on a Cloud have been discussed in previous work [120]. Sun et al. [119] provide an overview of the current state of the art on the development of customised analytics solutions on customers' premises and elaborate on some of the challenges to enable analytics and analytics as a service on the Cloud. Some of the potential business models proposed in their work include:

- **Hosting customer analytics jobs in a shared platform:** suitable for an enterprise or organisation that has multiple analytics departments. Traditionally, these departments have to develop their own analytics solutions and maintain their own clusters. With a shared platform they can upload their solutions to execute on a shared infrastructure, therefore reducing operation and maintenance costs. As discussed beforehand, techniques have been proposed for resource allocation and scheduling of Big Data analytics tasks on the Cloud [93,108].
- **A full stack designed to provide customers with end-to-end solutions:** appropriate for companies that do not have expertise on analysis. In this model, analytical service providers publish domain-specific analytical stream templates as services. The provider is responsible for hosting the software stack and managing the resources necessary to perform the analyses. Customers who subscribe to the services just need to upload their data, configure the templates, receive models, and perform the proper model scoring.
- **Expose analytics models as hosted services:** analytics capabilities are hosted on the Cloud and exposed to customers as

services. This model is proposed to companies that do not have enough data to make good predictions. Providers upload their models, which are consumed by customers via scoring services provided by the Cloud.

To make Big Data analytics solutions more affordable, Sun et al. [119] also propose cost-effective approaches that enable multi-tenancy at several levels. They discuss the technical challenges on isolating analytical artefacts. Hsueh et al. [73] discuss issues related to pricing and Service Level Agreements (SLAs) on a platform for personalisation in a wellness management platform built atop a Cloud infrastructure. Krishna and Varma [87] envision two types of services for Cloud analytics: (i) Analytics as a Service (AaaS), where analytics is provided to clients on demand and they can pick the solutions required for their purposes; and (ii) Model as a Service (MaaS) where models are offered as building blocks for analytics solutions.

Bhattacharya et al. [18] introduced IVOCA, a solution for providing managed analytics services for CRM. IVOCA provides functionalities that help analysts better explore data analysis tools to reduce the time to insight and improve the repeatability of CRM analytics. Also in the CRM realm, KXEN [89] offers a range of products for performing analytics, some of which can run on the Cloud. Cloud Prediction is a predictive analytics solution for Salesforce.com. With its Predictive Lead Scoring, Predictive Offers, and Churn Prediction, customers can leverage the CRM, mobile, and social data available in the Cloud to score leads based on which one can create sales opportunities; create offers that have a higher likelihood to be accepted based on a prediction of offers and promotions; and gain insights into which customers a company is at risk of losing.

Cloud-enabled Big Data analytics poses several challenges with respect to replicability of analyses. When not delivered by a Cloud, analytics solutions are customer-specific and models often have to be updated to consider new data. Cloud solutions for analytics need to balance generality and usefulness. Previous work also discusses the difficulty of replicating activities of text analytics [107]. An analytical pathway is proposed to link business objectives to an analytical flow, with the goal of establishing a methodology that illustrates and possibly supports repeatability of analytical processes when using complex analytics. King [84], whilst discussing some of the problems in buying predictive analytics, provides a best practice framework based on five steps, namely training, assessment, strategy, implementation, and iteration.

Chen et al. [34] envision an analytics ecosystem where data services aggregate, integrate, and provide access to public and private data by enabling partnerships among data providers, integrators, aggregators, and clients; these services are termed as DaaS. Atop DaaS, a range of analytics functionalities that explore the data services are offered to customers to boost productivity and create value. This layer is viewed as AaaS. Similar to the previously described work, they discuss a set of possible business models that range from proprietary, where both data and models are kept private, to co-developing models where both data and analytics models are shared among the parties involved in the development of the analytics strategy or services.

7. Other challenges

In business models where high-level analytics services may be delivered by the Cloud, human expertise cannot be easily replaced by machine learning and Big Data analysis [99]; in certain scenarios, there may be a need for human analysts to remain in the loop [91]. Management should adapt to Big Data scenarios and deal with challenges such as how to assist human analysts in gaining insights and how to explore methods that can help managers in making quicker decisions.

Application profiling is often necessary to estimate the costs of running analytics on a Cloud platform. Users need to develop their

applications to target Cloud platforms; an effort that should be carried out only after estimating the costs of transferring data to the Cloud, allocating virtual machines, and running the analysis. This cost estimation is not a trivial task to perform in current Cloud offerings. Although best practices for using some data processing services are available [49], there should be tools that assist customers to estimate the costs and risks of performing analytics on the Cloud.

Data ingestion by Cloud solutions is often a weak point, whereas debugging and validation of developed solutions is a challenging and tedious process. As discussed earlier, the manner analytics is executed on Cloud platforms resembles the batch job scenario: users submit a job and wait until tasks are executed and then download the results. Once an analysis is complete, they download sample results that are enough to validate the analysis task and after that perform further analysis. Current Cloud environments lack this interactive process, and techniques should be developed to facilitate interactivity and to include analysts in the loop by providing means to reduce their time to insight. Systems and techniques that iteratively refine answers to queries and give users more control of processing are desired [70].

Furthermore, market research shows that inadequate staffing and skills, lack of business support, and problems with analytics software are some of the barriers faced by corporations when performing analytics [112]. These issues can be exacerbated by the Cloud as the resources and analysts involved in certain analytics tasks may be offered by a Cloud provider and may move from one customer engagement to another. In addition, based on survey responses, currently most analytics updates and scores of methods occur daily to annually; which can become an issue for analytics on streaming data. Russom [112] also highlights the importance of advanced data visualisation techniques and advanced analytics – such as analysis of unstructured, large data sets and streams – to organisations in the next few years.

Chen et al. [32] foresee the emergence of what they termed as Business Intelligence and Analytics (BI&A) 3.0, which will require underlying mobile analytics and location and context-aware techniques for collecting, processing, analysing, and visualising large scale mobile and sensor data. Many of these tools are still to be developed. Moreover, moving to BI&A 3.0 will demand efforts on integrating data from multiple sources to be processed by Cloud resources, and using the Cloud to assist decisions by mobile device users.

More recently, terms such as Analytics as a Service (AaaS) and Big Data as a Service (BDaaS) are becoming popular. They comprise services for data analysis similarly as IaaS offers computing resources. However, these analytics services still lack well defined contracts since it may be difficult to measure quality and reliability of results and input data, provide promises on execution times, and guarantees on methods and experts responsible for analysing the data. Therefore, there are fundamental gaps on tools to assist service providers and clients to perform these tasks and facilitate the definition of contracts for both parties.

8. Summary and conclusions

The amount of data currently generated by the various activities of the society has never been so big, and is being generated in an ever increasing speed. This Big Data trend is being seen by industries as a way of obtaining advantage over their competitors: if one business is able to make sense of the information contained in the data reasonably quicker, it will be able to get more customers, increase the revenue per customer, optimise its operation, and reduce its costs. Nevertheless, Big Data analytics is still a challenging and time demanding task that requires expensive software, large computational infrastructure, and effort.

Cloud computing helps in alleviating these problems by providing resources on-demand with costs proportional to the actual usage. Furthermore, it enables infrastructures to be scaled up and down rapidly, adapting the system to the actual demand.

Although Cloud infrastructure offers such elastic capacity to supply computational resources on demand, the area of Cloud-supported analytics is still in its early days. In this paper, we discussed the key stages of analytics workflows, and surveyed the state-of-the-art of each stage in the context of Cloud-supported analytics. Surveyed work was classified in three key groups: Data Management (which encompasses data variety, data storage, data integration solutions, and data processing and resource management), Model Building and Scoring, and Visualisation and User Interactions. For each of these areas, ongoing work was analysed and key open challenges were discussed. This survey concluded with an analysis of business models for Cloud-assisted data analytics and other non-technical challenges.

The area of Big Data Computing using Cloud resources is moving fast, and after surveying the current solutions we identified some key lessons:

- There are plenty of solutions for Big Data related to Cloud computing. Such a large number of solutions have been created because of the wide range of analytics requirements, but they may, sometimes, overwhelm non-experienced users. Analytics can be descriptive, predictive, prescriptive; Big Data can have various levels of variety, velocity, volume, and veracity. Therefore, it is important to understand the requirements in order to choose appropriate Big Data tools;
- It is also clear that analytics is a complex process that demands people with expertise in cleaning up data, understanding and selecting proper methods, and analysing results. Tools are fundamental to help people perform these tasks. In addition, depending on the complexity and costs involved in carrying out these tasks, providers who offer Analytics as a Service or Big Data as a Service can be a promising alternative compared to performing these tasks in-house;
- Cloud computing plays a key role for Big Data; not only because it provides infrastructure and tools, but also because it is a business model that Big Data analytics can follow (e.g. Analytics as a Service (AaaS) or Big Data as a Service (BDaaS)). However, AaaS/BDaaS brings several challenges because the customer and provider's staff are much more involved in the loop than in traditional Cloud providers offering infrastructure/platform/software as a service.

Recurrent themes among the observed future work include (i) the development of standards and APIs enabling users to easily switch among solutions and (ii) the ability of getting the most of the elasticity capacity of the Cloud infrastructure. The latter includes expressive languages that enable users to describe the problem in simple terms whilst decomposing such high-level description in highly concurrent subtasks and keeping good performance efficiency even for large numbers of computing resources. If this can be achieved, the only limitations for an arbitrary short processing time would be market issues, namely the relation between the cost for running the analytics and the financial return brought for the obtained knowledge.

References

- [1] D.J. Abadi, Data management in the cloud: Limitations and opportunities, *IEEE Data Engineering Bulletin* 32 (1) (2009) 3–12.
- [2] Amazon redshift, <http://aws.amazon.com/redshift/>.
- [3] Amazon data pipeline, <http://aws.amazon.com/datapipeline/>.
- [4] Amazon Elastic MapReduce (EMR), <http://aws.amazon.com/elasticmapreduce/>.
- [5] Amazon Kinesis, <http://aws.amazon.com/kinesis/developer-resources/>.
- [6] R. Ananthanarayanan, K. Gupta, P. Pandey, H. Pucha, P. Sarkar, M. Shah, R. Tewari, Cloud Analytics: Do We Really Need to Reinvent the Storage Stack? in: *Proceedings of the Conference on Hot Topics in Cloud Computing (HotCloud 2009)*, USENIX Association, Berkeley, USA, 2009.
- [7] G. Andrienko, N. Andrienko, S. Wrobel, Visual analytics tools for analysis of movement data, *SIGKDD Explor. Newsl.* 9 (2) (2007) 38–46.
- [8] Announcing Suro: Backbone of Netflix's Data Pipeline, <http://techblog.netflix.com/2013/12/announcing-suro-backbone-of-netflix.html>.
- [9] Apache S4: distributed stream computing platform, <http://incubator.apache.org/s4/>.

- [10] Apache Hadoop, <http://hadoop.apache.org>.
- [11] Apache Mahout, <http://mahout.apache.org>.
- [12] Apache Samza, <http://samza.incubator.apache.org>.
- [13] M. Armbrust, A. Fox, R. Griffith, A.D. Joseph, R.H. Katz, A. Konwinski, G. Lee, D.A. Patterson, A. Rabkin, I. Stoica, M. Zaharia, Above the Clouds: A Berkeley View of Cloud Computing, Technical report UCB/Eecs-2009-28, Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, USA (February 2009).
- [14] Attention, shoppers: Store is tracking your cell, New York Times. URL <http://www.nytimes.com/2013/07/15/business/attention-shopper-stores-are-tracking-your-cell.html>.
- [15] A. Balmin, K. Beyer, V. Ercegovic, J.M.F. Ozcan, H. Pirahesh, E. Shekita, Y. Sismanis, S. Tata, Y. Tian, A platform for eXtreme Analytics, IBM J. Res. Dev. 57 (3–4) (2013) 4:1–4:11.
- [16] R.S. Barga, J. Ekanayake, W. Lu, Project Daytona: Data Analytics as a Cloud Service, in: A. Kementsietsidis, M. A. V. Salles (Eds.), Proceedings of the International Conference of Data Engineering (ICDE 2012), IEEE Computer Society, 2012, pp. 1317–1320.
- [17] G. Bell, T. Hey, A. Szalay, Beyond the Data Deluge, Science 323 (5919) (2009) 1297–1298.
- [18] I. Bhattacharya, S. Godbole, A. Gupta, A. Verma, J. Achtermann, K. English, Enabling Analysts in Managed Services for CRM Analytics, in: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2009), ACM, New York, USA, 2009, pp. 1077–1086.
- [19] bigdata@csail, <http://bigdata.csail.mit.edu/>.
- [20] 'Big Data' has Big Potential to Improve Americans' Lives, Increase Economic Opportunities, Committee on Science, Space and Technology (April 2013). URL <http://science.house.gov/press-release>.
- [21] Birst Inc., <http://www.birst.com>.
- [22] R. Bonney, J.L. Shirk, T.B. Phillips, A. Wiggins, H.L. Ballard, A.J. Miller-Rushing, J.K. Parrish, Next steps for citizen science, Science 343 (2014) 1436–1437.
- [23] D. Borthakur, J. Gray, J.S. Sarma, K. Muthukkaruppan, N. Spiegelberg, H. Kuang, K. Ranganathan, D. Molkov, A. Menon, S. Rash, R. Schmidt, A. Aiyer, Apache Hadoop Goes Realtime at Facebook, in: Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD 2011), ACM, New York, USA, 2011, pp. 1071–1080.
- [24] C. Bunch, N. Chohan, C. Krintz, J. Chohan, J. Kupferman, P. Lakhina, Y. Li, Y. Nomura, An Evaluation of Distributed Datastores Using the AppScale Cloud Platform, in: Proceedings of the 3rd IEEE International Conference on Cloud Computing (Cloud 2010), IEEE Computer Society, Washington, USA, 2010, pp. 305–312.
- [25] 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.
- [26] B. Calder, J. Wang, A. Ogus, N. Nilakantan, A. Skjolsvold, S. McKelvie, Y. Xu, S. Srivastav, J. Wu, H. Simitci, J. Haridas, C. Uddaraju, H. Khatri, A. Edwards, V. Bedekar, S. Mainali, R. Abbasi, A. Agarwal, M.H. Fahim, M.H. Ikram, D. Bhardwaj, S. Dayanand, A. Adusumilli, M. McNett, S. Sankaran, K. Manivannan, L. Rigas, Windows Azure Storage: A Highly Available Cloud Storage Service with Strong Consistency, in: Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP 2011), ACM, New York, NY, USA, 2011, pp. 143–157.
- [27] 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.
- [28] P.H. Carns, W.B.L. III, R.B. Ross, R. Thakur, PVFS: A parallel file system for linux clusters, in: Proceedings of the 4th Annual Linux Showcase and Conference (ALS 2000), USENIX, Atlanta, USA, 2000, pp. 317–328.
- [29] J. Chang, K.T. Lim, J. Byrne, L. Ramirez, P. Ranganathan, Workload diversity and dynamics in big data analytics: implications to system designers, in: Proceedings of the 2nd Workshop on Architectures and Systems for Big Data, ACM, New York, NY, USA, 2012, pp. 21–26.
- [30] Y. Chen, S. Alspaugh, D. Borthakur, R. Katz, Energy Efficiency for Large-Scale MapReduce Workloads with Significant Interactive Analysis, in: Proceedings of the 7th ACM European Conference on Computer Systems (EuroSys 2012), ACM, New York, USA, 2012, pp. 43–56.
- [31] Y. Chen, S. Alspaugh, R. Katz, Interactive Analytical Processing in Big Data Systems: A Cross-Industry Study of MapReduce Workloads, Proceedings of the VLDB Endowment 5 (12) (2012) 1802–1813.
- [32] H. Chen, R.H.L. Chiang, V.C. Storey, Business Intelligence and Analytics: From Big Data to Big Impact, MIS Quarterly 36 (4) (2012) 1165–1188.
- [33] Q. Chen, M. Hsu, H. Zeller, Experience in Continuous analytics as a Service (CaaS), in: Proceedings of the 14th International Conference on Extending Database Technology, ACM, New York, USA, 2011, pp. 509–514.
- [34] Y. Chen, J. Kreulen, M. Campbell, C. Abrams, Analytics ecosystem transformation: A force for business model innovation, in: Proceedings of the 2011 Annual SRII Global Conference (SRII 2011), IEEE Computer Society, Washington, USA, 2011, pp. 11–20.
- [35] K. Chen, H. Xu, F. Tian, S. Guo, CloudVista: Visual Cluster Exploration for Extreme Scale Data in the Cloud, in: Scientific and Statistical Database Management, Springer, 2011, pp. 332–350.
- [36] N. Chohan, A. Gupta, C. Bunch, K. Prakasam, Hybrid Cloud Support for Large Scale Analytics and Web Processing, in: Proceedings of the 3rd USENIX Conference on Web Application Development (WebApps 2012), Boston, USA, 2012.
- [37] J. Choo, H. Park, Customizing Computational Methods for Visual Analytics with Big Data, IEEE Computer Graphics and Applications 33 (4) (2013) 22–28.
- [38] Cloud9 Analytics, <http://www.cloud9analytics.com>.
- [39] J. Cohen, B. Dolan, M. Dunlap, J.M. Hellerstein, C. Welton, MAD skills: new analysis practices for big data, Proceedings of the VLDB Endow 2 (2) (2009) 1481–1492.
- [40] A. Cuzzocrea, I.-Y. Song, K.C. Davis, Analytics over large-scale multidimensional data: the big data revolution!, in: Proceedings of the ACM 14th international workshop on Data Warehousing and OLAP, ACM, New York, NY, USA, 2011, pp. 101–104.
- [41] DataDirect Cloud, <http://cloud.datadirect.com/> (2013).
- [42] T.H. Davenport, J.G. Harris, Competing on Analytics: The New Science of Winning, Harvard Business Review Press, 2007.
- [43] T.H. Davenport, J.G. Harris, R. Morison, Analytics at Work: Smarter Decisions, Better Results, Harvard Business Review Press, 2010.
- [44] J. Davey, F. Mansmann, J. Kohlhammer, D. Keim, The future internet, Springer-Verlag, Berlin, Heidelberg, 2012, Ch. Visual Analytics: Towards Intelligent Interactive Internet and Security Solutions, pp. 93–104.
- [45] J. Dean, S. Ghemawat, MapReduce: Simplified Data Processing on Large Clusters, Communications of the ACM 51(1).
- [46] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, W. Vogels, Dynamo: Amazon's Highly Available Key-Value Store, SIGOPS Operating Systems Review 41 (6) (2007) 205–220.
- [47] E. Deelman, A. Chervenak, Data management challenges of data-intensive scientific workflows, in: Proceedings of the 8th IEEE International Symposium on Cluster Computing and the Grid (CCGrid'08), IEEE Computer Society, 2008, pp. 687–692.
- [48] P. Deepak, P.M. Deshpande, K. Murthy, Configurable and Extensible Multi-flows for Providing Analytics as a Service on the Cloud, in: Proceedings of the 2012 Annual SRII Global Conference (SRII 2012), 2012, pp. 1–10.
- [49] P. Deyhim, Best practices for Amazon EMR, White paper, Amazon (2013). URL http://media.amazonwebservices.com/AWS_Amazon_EMR_Best_Practices.pdf.
- [50] U. Fayyad, G. Piatetsky-Shapiro, P. Smyth, The KDD process for extracting useful knowledge from volumes of data, Commun. ACM 39 (11) (1996) 27–34.
- [51] R. Feldman, Techniques and applications for sentiment analysis, Communications of the ACM 56 (4) (2013) 82–89.
- [52] D. Fisher, R. DeLine, M. Czerwinski, S. Drucker, Interactions with Big Data Analytics, Interactions 19 (3) (2012) 50–59.
- [53] D. Fisher, I. Popov, S.M. Drucker, M. Schraefel, Trust me, I'm partially right: Incremental visualization lets analysts explore large datasets faster, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2012), ACM, New York, USA, 2012, pp. 1673–1682.
- [54] G.C. Fox, Large Scale Data Analytics on Clouds, in: Proceedings of the Fourth International Workshop on Cloud Data Management (CloudDB 2012), ACM, 2012, pp. 21–24.
- [55] B. Franks, Taming The Big Data Tidal Wave: Finding Opportunities in Huge Data Streams with Advanced Analytics, first ed., in: Wiley and SAS Business Series, Wiley, 2012.
- [56] FusionCharts, <http://www.fusioncharts.com/>.
- [57] S. Ghemawat, H. Gobioff, S.-T. Leung, The google file system, in: Proceedings of the 9th ACM Symposium on Operating Systems Principles (SOSP 2003), ACM, New York, USA, 2003, pp. 29–43.
- [58] Google App Engine, <http://developers.google.com/appengine/>.
- [59] Google Prediction API, <https://developers.google.com/prediction/>.
- [60] Google Analytics, <http://www.google.com/analytics/>.
- [61] Gooddata, <http://www.gooddata.com> (2013).
- [62] K. Goodhope, J. Koshy, J. Kreps, N. Narkhede, R. Park, J. Rao, V.Y. Ye, Building LinkedIn's Real-time Activity Data Pipeline, Bulletin of the Technical Committee on Data Engineering 35 (2) (2012) 33–45.
- [63] R.L. Grossman, What is analytic infrastructure and why should you care? ACM SIGKDD Explorations Newsletter 11 (1) (2009) 5–9.
- [64] A. Guazzelli, K. Stathatos, M. Zeller, Efficient Deployment of Predictive Analytics Through Open Standards and Cloud Computing, ACM SIGKDD Explorations Newsletter 11 (1) (2009) 32–38.
- [65] A. Guazzelli, M. Zeller, W.-C. Lin, G. Williams, PMML: An Open Standard for Sharing Models, The R Journal 1 (1) (2009) 60–65.
- [66] Z. Guo, G. Fox, Improving MapReduce Performance in Heterogeneous Network Environments and Resource Utilization, in: Proceedings of the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2012), IEEE, 2012, pp. 714–716.
- [67] D. Habich, W. Lehner, S. Richly, U. Alßmann, Using Cloud Technologies to Optimize Data-Intensive Service Applications, in: Proceedings of the IEEE 3rd International Conference on Cloud Computing (CLOUD 2010), 2010, pp. 19–26.
- [68] J. Han, H. E. G. Le, J. Du, Survey on NoSQL database, in: 6th International Conference on Pervasive Computing and Applications (ICPCA 2011), IEEE, Port Elizabeth, South Africa, 2011, pp. 363–366.
- [69] R. Hecht, S. Jablonski, NoSQL Evaluation—A Use Case Oriented Survey, in: Proceedings of the International Conference on Cloud and Service Computing (CSC 2011), IEEE, Hong Kong, 2011, pp. 336–341.
- [70] J.M. Hellerstein, R.A.A. Chou, C. Hidber, C. Olston, V. Raman, T. Roth, P.J. Haas, Interactive Data Analysis: the Control Project, Computer 32 (8) (1999) 51–59.
- [71] H. Herodotou, S. Babu, Profiling, What-if Analysis, and Cost-based Optimization of MapReduce Programs, Proceedings of the VLDB Endowment 4 (11) (2011) 1111–1122.

- [72] H. Herodotou, H. Lim, G. Luo, N. Borisov, L. Dong, F.B. Cetin, S. Babu, Starfish: A Self-tuning System for Big Data Analytics, in: Proceedings of the 5th Biennial Conference on Innovative Data Systems Research (CIDR 2011), 2011, pp. 261–272.
- [73] P.-Y. S. Hsueh, R.J. Lin, M.J. Hsiao, L. Zeng, S. Ramakrishnan, H. Chang, Cloud-based Platform for Personalization in a Wellness Management Ecosystem: Why, What, and How, in: Proceedings of the 6th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2010), 2010, pp. 1–8.
- [74] B. Huang, S. Babu, J. Yang, Cumulon: Optimizing Statistical Data Analysis in the Cloud, in: Proceedings of the ACM SIGMOD International Conference on Management of Data SIGMOD 2013, ACM, New York, USA, 2013, pp. 1–12.
- [75] IBM InfoSphere Streams, <http://www.ibm.com/software/products/en/infosphere-streams>.
- [76] IBM SmartCloud Enterprise, <http://www-935.ibm.com/services/us/en/cloud-enterprise/> (2012).
- [77] Infochimps cloud overview, <http://www.infochimps.com/infochimps-cloud/overview/>.
- [78] A. Iosup, A. Lascateu, N. Tapus, CAMEO: Enabling social networks for Massively Multiplayer Online Games through Continuous Analytics and Cloud Computing, in: Proceedings of the 9th Annual Workshop on Network and Systems Support for Games (NetGames 2010), 2010, pp. 1–6.
- [79] D. Jensen, K. Konkol, A. Mohindra, F. Naccarati, E. Sam, Business Analytics in the Cloud, White paper IBW03004-USEN-00, IBM (April 2012).
- [80] G. Jung, N. Gnanasambandam, T. Mukherjee, Synchronous Parallel Processing of Big-Data Analytics Services to Optimize Performance in Federated Clouds, in: Proceedings of the IEEE 5th International Conference on Cloud Computing (Cloud 2012), 2012, pp. 811–818.
- [81] H. Kasim, T. Hung, E.F.T. Legara, K.K. Lee, X. Li, B.-S. Lee, V. Selvam, S. Lu, L. Wang, C. Monterola, V. Jayaraman, Scalable Complex System Modeling for Sustainable City, in: The 6th IEEE International Scalable Computing Challenge (SCALE 2013) in conjunction with The 13th International Symposium on Cluster, Cloud and the Grid (CCGrid 2013), 2013.
- [82] D.S. Katz, S. Jha, M. Parashar, O. Rana, J.B. Weissman, Survey and Analysis of Production Distributed Computing Infrastructures, CoRR abs/1208.2649.
- [83] H. Kim, S. Chaudhari, M. Parashar, C. Marty, Online Risk Analytics on the Cloud, in: Proceedings of the 9th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2009), IEEE Computer Society, Washington, USA, 2009, pp. 484–489.
- [84] E.A. King, How to buy data mining: A framework for avoiding costly project pitfalls in predictive analytics, DMReview 15(10).
- [85] J. Kobielius, In-Database Analytics: The Heart of the Predictive Enterprise, Technical report, Forrester Research, Inc., Cambridge, USA (Nov. 2009).
- [86] R. Kosara, J. Mackinlay, Storytelling: The Next Step for Visualization, Computer 46 (5) (2013) 44–50.
- [87] P.R. Krishna, K.I. Varma, Cloud Analytics: A Path Towards Next Generation Affordable BI, White paper, Infosys (2012).
- [88] A. Kumar, F. Niu, C. Ré, Hazy: Making it Easier to Build and Maintain Big-Data Analytics, Communications of the ACM 56 (3) (2013) 40–49.
- [89] KXEN, <http://www.kxen.com>.
- [90] J.K. Laurila, D. Gatica-Perez, I. Aad, J. Blom, O. Bornet, T.-M.-T. Do, O. Dousse, J. Eberle, M. Miettinen, The Mobile Data Challenge: Big Data for Mobile Computing Research (2012). URL: http://research.nokia.com/files/public/MDC2012_Overview_LaurilaGaticaPerezEtAl.pdf.
- [91] D. Lazer, R. Kennedy, G. King, A. Vespignani, The Parable of google flu: Traps in big data analysis, Science 343 (2014) 1203–1205.
- [92] N. Leavitt, Will NoSQL Databases Live Up to Their Promise? Computer 43 (2) (2010) 12–14.
- [93] G. Lee, B.-G. Chun, R.H. Katz, Heterogeneity-Aware Resource Allocation and Scheduling in the Cloud, in: Proceedings of the 3rd USENIX conference on Hot topics in Cloud computing (HotCloud 2011), USENIX Association, Berkeley, USA, 2011.
- [94] K.-H. Lee, Y.-J. Lee, H. Choi, Y.D. Chung, B. Moon, Parallel Data Processing with MapReduce: A Survey, SIGMOD Record 40 (4) (2011) 11–20.
- [95] T.-Y. Lee, X. Tong, H.-W. Shen, P.C. Wong, S. Hagos, L.R. Leung, Feature Tracking and Visualization of the Madden-Julian Oscillation in Climate Simulation, IEEE Computer Graphics and Applications 33 (4) (2013) 29–37.
- [96] Q. Liao, L. Shi, C. Wang, Visual analysis of large-scale network anomalies, IBM J. Res. Dev. 57 (3/4) (2013) 13:1–13:12.
- [97] X. Li, R.N. Calheiros, S. Lu, L. Wang, H. Palit, Q. Zheng, R. Buyya, Design and Development of an Adaptive Workflow-Enabled Spatial-Temporal Analytics Framework, in: Proceedings of the IEEE 18th International Conference on Parallel and Distributed Systems (ICPADS 2012), IEEE Computer Society, Singapore, 2012, pp. 862–867.
- [98] S. Lu, R.M. Li, W.C. Tjhi, K.K. Lee, L. Wang, X. Li, D. Ma, A framework for cloud-based large-scale data analytics and visualization: Case study on multiscale climate data, in: Proceedings of the IEEE 3rd International Conference on Cloud Computing Technology and Science (CloudCom 2011), IEEE Computer Society, Washington, USA, 2011, pp. 618–622.
- [99] A. McAfee, E. Brynjolfsson, Big data: The management revolution, Harv. Bus. Rev. (2012) 60–68.
- [100] S. Melnik, A. Gubarev, J.J. Long, G. Romer, S. Shivakumar, M. Tolton, T. Vassilakis, Dremel: Interactive Analysis of Web-Scale Datasets, Proceedings of the VLDB Endowment 3 (1–2) (2010) 330–339.
- [101] T. Menzies, T. Zimmermann, Software analytics: So what? IEEE Software 30 (4) (2013) 31–37.
- [102] L. Moreau, P. Groth, S. Miles, J. Vazquez-Salceda, J. Ibbotson, S. Jiang, S. Munroe, O. Rana, A. Schreiber, V. Tan, L. Varga, The provenance of electronic data, Communications of the ACM 51 (4) (2008) 52–58.
- [103] S. Murray, Interactive Data Visualization for the Web, O'Reilly Media, 2013.
- [104] panXpan, <https://www.panxpan.com>.
- [105] PivotLink AnalyticsCLOUD, <http://www.pivotlink.com/products/analyticscloud>.
- [106] Prime Minister joins Sir Ka-shing Li for launch of £90m initiative in big data and drug discovery at Oxford, <http://www.cs.ox.ac.uk/news/639-full.html> (May 2013).
- [107] L. Proctor, C.A. Kieliszewski, A. Hochstein, S. Spangler, Analytical pathway methodology: Simplifying business intelligence consulting, in: Proceedings of the Annual SRII Global Conference (SRII 2011), 2011, pp. 495–500.
- [108] M. Rahman, X. Li, H. Palit, Hybrid Heuristic for Scheduling Data Analytics Workflow Applications in Hybrid Cloud Environment, in: Proceedings of the IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum (IPDPSW), 2011, pp. 966–974.
- [109] K. Reda, A. Febretti, A. Knoll, J. Aurisano, J. Leigh, A. Johnson, M.E. Papka, M. Hereld, Visualizing Large, Heterogeneous Data in Hybrid-Reality Environments, IEEE Computer Graphics and Applications 33 (4) (2013) 38–48.
- [110] T.C. Redman, Data Quality for the Information Age, Artech House, 1997.
- [111] Right90, <http://www.right90.com>.
- [112] P. Russom, Big Data Analytics, TDWI best practices report, The Data Warehousing Institute (TDWI) Research (2011).
- [113] S. Sakr, A. Liu, D. Batista, M. Alomari, A survey of large scale data management approaches in cloud environments, IEEE Communications Surveys Tutorials 13 (3) (2011) 311–336.
- [114] Salesforce, <http://www.salesforce.com>.
- [115] SAP HANA One, <http://www.saphana.com/community/solutions/cloud-info> (2013).
- [116] SAP Crystal Solutions, <http://www.crystalreports.com/>.
- [117] F. Schmuck, R. Haskin, GPPS: A Shared-Disk File System for Large Computing Clusters, in: Proceedings of the 1st Conference on File and Storage Technologies (FAST'02), Monterey, USA, 2002, pp. 231–244.
- [118] F. Schom, F. Stahl, G. Vossen, Marketplaces for data: An initial survey, SIGMOD Record 42 (1) (2013) 15–26.
- [119] X. Sun, B. Gao, L. Fan, W. An, A Cost-Effective Approach to Delivering Analytics as a Service, in: Proceedings of the 19th IEEE International Conference on Web Services (ICWS 2012), Honolulu, USA, 2012, pp. 512–519.
- [120] X. Sun, B. Gao, Y. Zhang, W. An, H. Cao, C. Guo, W. Sun, Towards delivering analytical solutions in cloud: Business models and technical challenges, in: Proceedings of the IEEE 8th International Conference on e-Business Engineering (ICEBE 2011), IEEE Computer Society, Washington, USA, 2011, pp. 347–351.
- [121] Storm: distributed and fault-tolerant realtime computation, <http://storm.incubator.apache.org>.
- [122] The Intel science and technology center for big data, <http://istc-bigdata.org>.
- [123] I. T. Tabor Communications, The UberCloud HPC Experiment: Compendium of Case Studies, Tech. rep. (2013).
- [124] W. Tantisiriroj, S.W. Son, S. Patil, S.J. Lang, G. Gibson, R.B. Ross, On the duality of data-intensive file system design: reconciling HDFS and PVFS, in: Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC 2011), ACM, New York, NY, USA, 2011, pp. 67:1–67:12.
- [125] A. Thusoo, Z. Shao, S. Anthony, D. Borthakur, N. Jain, J.S. Sarma, R. Murthy, H. Liu, Data warehousing and analytics infrastructure at Facebook, in: Proceedings of the 2010 International Conference on Management of Data, ACM, New York, NY, USA, 2010, pp. 1013–1020.
- [126] S.R. Upadhyaya, Parallel Approaches to Machine Learning—A Comprehensive Survey, Journal of Parallel Distributed Computing 73 (3) (2013) 284–292.
- [127] Unlocking Game-Changing Wireless Capabilities: Cisco and SITA help Copenhagen Airport Develop New Services for Transforming the Passenger Experience, Customer case study, CISCO (2012). URL http://www.cisco.com/en/US/prod/collateral/wireless/c36_696714_00_copenhagen_airport_cs.pdf.
- [128] S. Venugopal, R. Buyya, K. Ramamohanarao, A taxonomy of data grids for distributed data sharing, management and processing, ACM Comput. Surv. 38 (1) (2006) 1–53.
- [129] F.B. Viegas, M. Wattenberg, F. van Ham, J. Kriss, M. McKeon, ManyEyes: A Site for Visualization at Internet Scale, IEEE Trans. Vis. Comput. Graphics 13 (6) (2007) 1121–1128.
- [130] H. Wang, Integrity Verification of Cloud-Hosted Data Analytics Computations, in: Proceedings of the 1st International Workshop on Cloud Intelligence (Cloud-I 2012), ACM, New York, NY, USA, 2012.
- [131] C. Wang, K. Schwan, V. Talwar, G. Eisenhauer, L. Hu, M. Wolf, A Flexible Architecture Integrating Monitoring and Analytics for Managing Large-Scale Data Centers, in: Proceedings of the 8th ACM International Conference on Autonomic Computing (ICAC 2011), ACM, New York, USA, 2011, pp. 141–150.
- [132] Windows Azure HDInsight, <http://www.windowsazure.com/en-us/documentation/services/hdinsight/>.
- [133] I.H. Witten, E. Frank, M.A. Hall, Data Mining: Practical Machine Learning Tools and Techniques, third ed., Morgan Kaufmann, 2011.
- [134] XSEDE, <http://www.xsede.org/>.
- [135] H. Xu, Z. Li, S. Guo, K. Chen, CloudVista: Interactive and Economical Visual Cluster Analysis for Big Data in the Cloud, Proceedings of the VLDB Endowment 5 (12) (2012) 1886–1889.
- [136] P.S. Yu, On mining big data, in: J. Wang, H. Xiong, Y. Ishikawa, J. Xu, J. Zhou (Eds.), Web-Age Information Management, in: Lecture Notes in Computer Science, vol. 7923, Springer-Verlag, Berlin, Heidelberg, 2013, p. XIV.

- [137] G. Yunwen, W. Shaochun, Y. Bowen, L. Jiazheng, The Methods of Data Prefetching Based on User Model in Cloud Computing, in: Proceedings of the 2011 International Conference on Internet of Things (ITHINGSCPCOM 2011), IEEE Computer Society, Washington, DC, USA, 2011, pp. 463–466.
- [138] Zementis – adaptive decision technology, <http://www.zementis.com> (2012).
- [139] X. Zhang, E. Zhang, B. Song, F. Wei, Towards Building an Integrated Information Platform for Eco-city, in: Proceedings of the 7th International Conference on e-Business Engineering (ICEBE 2010), 2010, pp. 393–398.
- [140] P. Zikopoulos, C. Eaton, P. Zikopoulos, Understanding Big Data: Analytics for Enterprise Class Hadoop and Streaming Data, McGraw-Hill Companies, Inc, 2012.



Silvia Bianchi is a Research Staff Member in the Service Systems group of IBM Research Brazil. She joined IBM in March 2012. Silvia received B.Sc. degree in Computer Science from the Federal University of Santa Catarina (UFSC), Brazil, M.Sc. degree in Computer Science from Paul Sabatier University (UPS), France, and Ph.D. in Computer Science from the University of Neuchatel (Unine) in Switzerland. She is currently involved in projects on Cloud Computing, Peer-to-Peer and Publish/Subscribe.



Marcos Dias de Assuncao, a former member of the research staff at IBM, is interested in workload migration, resource management in Cloud computing, and techniques for big data analysis. Marcos obtained Ph.D. in Computer Science and Software Engineering (2009) from the University of Melbourne, Australia.



Marco A.S. Netto is a Researcher at IBM Research Brazil, where he works on Cloud Computing and Analytics related projects. Marco obtained his Ph.D. in Computer Science and Software Engineering (2010) from the University of Melbourne, Australia. His research interests are Cluster/Grid/Cloud Computing with focus on SLA management, virtualisation, performance evaluation, job scheduling, Quality-of-Service, and optimisation issues.



Dr. Rodrigo N. Calheiros is a Research Fellow in the Department of Computing and Information Systems, the University of Melbourne, Australia. Since 2010, he is a member of the CLOUDS Lab of the University of Melbourne, where he researches various aspects of cloud computing. He works in the field of Cloud computing since 2008. His research interests also include virtualization, grid computing, and simulation and emulation of distributed systems.



Dr. Rajkumar Buyya is 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 the founding CEO of Manjrasoft, a spin-off company of the University, commercialising its innovations in Cloud Computing. He has authored 400 publications and four text books. He is one of the highly cited authors in computer science and software engineering worldwide.