Cloud computing in e-science: research challenges and opportunities

Article

Yang, X., Wallom, D., Waddington, S., Wang, J., Shaon, A., Matthews, B., Wilson, M., Guo, Y., Guo, L., Blower, J. D., Athanasios, V. V., Liu, K. and Kershaw, P. (2014) Cloud computing in e-science: research challenges and opportunities. *Journal of Supercomputing*, 70 (1). pp. 408-464. ISSN 0920-8542 doi: <https://doi.org/10.1007/s11227-014-1251-5> Available at <https://centaur.reading.ac.uk/36995/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1007/s11227-014-1251-5>

Publisher: Springer

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Cloud Computing in e-Science: Research Challenges and Opportunities

Xiaoyu Yang^{1,10}, David Wallom², Simon Waddington³, Jianwu Wang⁴, Arif Shaon⁵, Brian Matthews⁵, Michael Wilson⁵, Yike Guo⁶, Li Guo⁶, Jon Blower¹, Athanasios V. Vasilakos⁷, Kecheng Liu⁸, Philip Kershaw⁹

1. *e-Science Research Centre, University of Reading, UK*
2. *Oxford e-Research Centre, University of Oxford, UK*
3. *Centre for e-Research, King's College London, UK*
4. *San Diego Supercomputer Centre, University of California, San Diego, USA*
5. *e-Science Department, Rutherford Appleton Laboratory, STFC, UK*
6. *Department of Computing, Imperial College London, UK*
7. *Department of Computer and Telecommunications Engineering, University of Western Macedonia, Greece*
8. *Informatics Research Centre, University of Reading, UK*
9. *NCEO/Centre for Environmental Data Archival, Rutherford Appleton Laboratory, STFC, UK*
10. *Computing Network Information Centre, Chinese Academy of Sciences, Beijing, China*

Abstract

SOA (Service Oriented Architecture), workflow, the Semantic Web, and Grid computing are key enabling information technologies in the development of increasingly sophisticated e-Science infrastructures and application platforms. While the emergence of Cloud computing as a new computing paradigm has provided new directions and opportunities for e-Science infrastructure development, it also presents some challenges. Scientific research is increasingly finding that it is difficult to handle “big data” using traditional data processing techniques. Such challenges demonstrate the need for a comprehensive analysis on using the above mentioned informatics techniques to develop appropriate e-Science infrastructure and platforms in the context of Cloud computing. This survey paper describes recent research advances in applying informatics techniques to facilitate scientific research particularly from the Cloud computing perspective. Our particular contributions include identifying associated research challenges and opportunities, presenting lessons learned, and describing our future vision for applying Cloud computing to e-Science. We believe our research findings can help indicate the future trend of e-Science, and can inform funding and research directions in how to more appropriately employ computing technologies in scientific research. We point out the open research issues hoping to spark new development and innovation in the e-Science field.

Keywords: e-Science, e-Research, Informatics, Cloud computing, Semantic Web, Grid computing, Workflow, Digital research, Big data

1 Introduction

Next generation scientific research has radically changed the way in which science is carried out [1], [2]. With the assistance of modern e-Infrastructure that integrates high performance computing, large capacity data storage facilities and high speed network infrastructure, the exploration of previously unknown problems can now be solved by simulation, generation and analysis of large amounts of data, sharing of geographically distributed resources (e.g. computing facilities, data, scripts, experimental plans, workflows) and global research collaboration.

e-Science facilitates new dimensions of research and experimentation through global interdisciplinary collaboration involving both people and shared resources. These mega-scale collaborations are underpinned by the e-Science infrastructures that support and enhance the scientific process by enabling more efficient production, analysis and sharing of experiments, results and other related information [3]. The capability and sophistication of the e-Science infrastructures are steadily improving, driven by the exponential increases in computing power and high speed networks. This is resulting in a dramatic rise in the volume of data generated and published by various e-Science related activities. e-Science is now in an age of “big data”. For example, the e-Science data generated from different areas, such as sensors,

satellites and high-performance computer simulations have been measured in the excess of terabytes every year and are expected to inflate significantly over the next decade. The beamlines of the DIAMOND Synchrotron¹ facility [4] in the UK are now used by several hundred scientists every year, generating several terabytes of data per day and accumulating hundreds of terabytes of data per year. Automatic, reproducible, reusable and repeatable research involves the automatic coordination of various tasks and data involved in a research study, and provenance of the research results. All these require the employment of Information and Communication Technologies (ICT) to build appropriate e-infrastructure that can encompass the computation facility, data storage, networks, software, people, and training in a holistic manner [2].

Features of next generation scientific research, the e-Science requirements imposed by these features, and major enabling technologies are summarised in Figure 1.

Enabling technologies	e-Science requirements	New scientific research features
Workflow, Semantic Web, etc.	Repeatability, Reusability, Intelligent, ...	Automatic, Reproducible,
Semantic Web, Hadoop, SaaS, PaaS, IaaS, DaaS, etc.	Data mining, Data publishing Data curation, Linked data ...	Data intensive, Big data
SOA, Grid, Virtualisation, Web 2.0, SaaS, PaaS, IaaS, etc.	Integrated platforms, Virtual lab, VO ...	Research collaboration
Grid, Semantic Web, SaaS, PaaS, IaaS, NaaS, etc.	Open data, Knowledge base Intelligence collection, ...	Resource sharing

Figure 1 A summary of e-Science requirements and key enabling technologies

As shown in the Figure 1, Web services and SOA, workflow, Semantic Web, Grid computing and Cloud computing (e.g. SaaS, PaaS, IaaS), etc. are some of key enabling digital technologies for developing appropriate e-Infrastructure and application-oriented platforms. In particular, Cloud computing can be used in e-Science to provide scalable IT infrastructure, QoS-assured services and a customisable computing environment. As a new computing paradigm, it imposes many research challenges and opportunities involving technical, cultural and business issues. In order to report on the recent advances in applying these information technologies to e-Science, we have produced this survey paper, discussing the use of these technologies and identifying research challenges and opportunities.

The paper is structured as follows: Section 2 discusses some research challenges and opportunities of applying Web services and SOA in e-Science from the perspective of service computing. Section 3 presents research aspects and community practices for using workflow in e-Science, and identifies some research challenges and opportunities of using workflow in scientific research. Issues of big data processing in the Cloud are also discussed in this section. Section 4 gives an overview of the role of the Semantic Web in e-Science and identifies some key challenges in its application to this domain. Section 5 discusses the evolution from Grid computing to Cloud computing by focusing on the comparison of essential characteristics of

¹ The Diamond Light Source - <http://www.diamond.ac.uk>

both, and presents a vision on how e-Science should evolve in the Cloud computing era, known as “e-Science 2.0”. A layered e-Science 2.0 framework is proposed. Section 6 discusses research opportunities and challenges in Cloud storage, as it opens up research opportunities for data intensive research, which have not traditionally been supported by custom e-Science infrastructure such as Grid computing. Section 7 discusses some Cloud computing issues in e-Science from high performance computing perspective, with some lessons learned and future visions.

2 Web Service and SOA in e-Science

Web services and Service-Oriented Architecture (SOA) are important distributed computing technologies that are currently widely used. This section describes issues and challenges in employing them to support the scientific research from the Cloud perspective.

2.1 From e-Research Models to SOA and Service Computing

Scientific resources include computing resources, storage resources, data, legacy code and scripts, sensors, instruments, etc. which are usually distributed across administrative domains. Providing effective access to them without compromising the security and local control is a fundamental issue in e-Science. An e-Research model [5] has been proposed to accommodate this need where “resources” can be accessed through “services”. This is an approach whereby resource owners make their resources available to collaborating researchers by providing a well-defined interface specifying the operations that can be performed on, or with, a given resource, e.g. submitting a compute job or accessing a set of data. These services can often be accessed through a Web-based or machine-oriented interface, allowing both their interactive use and their integration with other services to provide higher-level functionality.

The term ‘Service Computing’ has emerged to encompass the technologies of Web services, Service Oriented Architecture (SOA) and workflow. Service computing can be defined as a multi-disciplinary domain that covers the science and technology of bridging the gap between Business Services and IT Services [6]. It aims to enable IT services and computing technology to provide business services more efficiently and effectively. The supporting technology suite includes for example Web services, SOA, and business process integration and management. (In e-Science, “business processes” are usually referred to as scientific workflows). Although the idea of service computing is business-oriented, the concept has been widely used in e-Science as a basis for modern cyber-infrastructure. For example, in the UK TSB funded MaterialsGrid² [7], [8] project, the existing scientific code and scripts were wrapped as Web services / Grid services and integrated as a service-oriented workflow to calculate the properties of a material (Figure 2). In the European NETMAR [9] project, Web services and SOA are employed to develop a pilot European Marine Information System (EUMIS) for searching, downloading and integrating satellite, in-situ and model data from ocean and coastal areas. The Global Earth Observation System of Systems (GEOSS) is a global system from which users of earth observations can use GEOSS to search and use the distributed data, information, tools and services. GEOSS does not aim to develop a new system; instead it employs a System of Systems approach to link existing systems globally, and thus deliver value-added functionalities. The GEOSS Common Infrastructure (GCI) is entirely based on SOA principles [10].

² <http://www.materialsgrid.org/>

Figure 2 A service-oriented workflow in MaterialsGrid used to simulate elastic constant of a material using CASTEP quantum mechanical simulation code. Each block is wrapped as a reusable Web service and integrated as a service-oriented workflow. In [11], this is typed as a sophisticated simulation which combines single simulations, two- consecutive simulations, and parameter sweep.

The emergence of Cloud computing has resulted in a new computing paradigm for e-Science, where the researcher or research institute does not have to maintain physical infrastructure (e.g. high performance computing resources, data storage resources) but instead purchases infrastructure services from dedicated providers. This computing paradigm provides many potential benefits including a scalable IT infrastructure, QoS-assured services and a customisable computing environment. However, it also presents many research challenges and opportunities.

2.2 Research Challenges and Opportunities of Service Computing

Challenges and opportunities of providing service computing using the Cloud include how to enable: (i) Quality-aware service delivery, (ii) Service monitoring, and (iii) Service metering?

2.2.1 Service Level Agreement

Web services have now been widely used in e-Science, service-oriented infrastructure, and various computing paradigms (e.g. Cloud computing, Grid computing, service computing). Thus, ensuring Quality of Service (QoS) is becoming an increasingly important topic. One approach to achieve this is to employ Service Level Agreements (SLAs) to serve as a bilateral contract that exists between a customer and a service provider to specify the user requirements, quality of service, responsibilities and obligations³. SLAs can contain numerous service performance metrics with corresponding Service Level Objectives (SLOs). An SLA

³ <http://www.gridipedia.eu>

describes quality of service and other commitments by a service provider to meet the obligations (e.g. in exchange for financial commitments based on an agreed schedule of prices and payment⁴).

In the Cloud-based computing paradigm, while the relationship between a researcher / scientist and service provider is inherently a “customer – service provider” relationship, the service provider and infrastructure provider also have to establish a similar relationship. As the service provider has both customer side and infrastructure provider side commitments, employing an appropriate SLA to guarantee the delivery of QoS-assured service becomes critical.

While the concept of SLA is simple, the underlying supporting infrastructure can be complex. For example, if there are several service providers that can provide similar functional services to researchers, how can service providers publish their services? How can researchers search those services not only by functionality, but also by QoS requirements? One possible solution is to use a Service Marketplace that can provide a store (or “service registry”) for published services. The service marketplace should return a list of Endpoint References (EPRs) of matching services with QoS constraints or price. The returned list needs to be ranked by QoS constraints or price. The service marketplace can be implemented using standard-based technologies (e.g. such as UDDI⁵, ebXML⁶). These aspects of SOA present research opportunities for e-Science in the Cloud computing environment

SLA negotiation is another important topic for SLA supporting infrastructure to ensure quality of service. The associated Cloud computing environment needs to provide facilities where researchers can negotiate with the service provider for service level requirements, in a certain granularity, by making an SLA proposal. The service provider then needs to check the resources to make sure whether the proposed service level requirements can be guaranteed, and decide whether to accept or reject the proposal. Once the SLA proposal is agreed by both parties, it becomes an SLA instance that serves as a contract for the researcher and the service provider throughout the whole session

Enforcing mechanisms are also important in SLA supporting infrastructure to ensure that terms defined in SLA can be guaranteed [12]. For example, an appropriate SLA enforcing component needs to contain the following sub-modules. (1) A Notifier, which can send notification messages to the service provider or customer when violation events occur. (2) A Rule Engine: based on the feature of the event, the Rule Engine can (i) retrieve relevant the policy from the Policy Store, (ii) get associated parameters from SLA instance, and/ or (iii) get domain knowledge from a knowledge base, to decide the recovery actions or rescheduling. (3) A Trigger, which is responsible for taking recovery actions or rescheduling actions defined by the Rule Engine. For example, these actions can include getting a new service deployed, or getting a new session created.

Providing a QoS assured services to customers (e.g. researchers) with minimal resource consumption cost and achieving customer satisfaction, whilst also guaranteeing the maximisation of the scientific goal or business objectives (e.g. margin profit) to the service provider and infrastructure provider within certain constraints, also presents challenges. For example, the resource consumption of running an application can be subject to several factors: (i) application workload feature, (ii) user interaction, (iii) network features, and (iv) mean time to failure. How can we use optimisation technology to find a resource with an associated

⁴ <http://www.gria.org/about-gria/a-business-perspective>

⁵ <http://www.oasis-open.org/committees/uddi-spec/>

⁶ http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=regrep

configuration that can guarantee service behaviour within the constraints and can maximise an objective function? How can we use a set of models (e.g. user behaviour modelling, resource behaviour modelling, uninterrupted fault-free application behaviour modelling, interactive application performance estimation) to describe aspects of the user, the resource and the application behaviour, which can then be combined to determine the behaviour of the service as a whole? Knowing this can help us better understand the features of the service so that we can estimate the service cost and reserve the appropriate resources in advance to ensure trade-off among QoS guarantees, cost, and margin profit to stakeholders in the value chain involving service customer, service provider and infrastructure provider.

2.2.2 Service Monitoring

In order to ensure the quality of service (e.g. of a scientific computing service) and to enable the elastic feature of Cloud computing and workload management [13], service monitoring is also required and can present research opportunities for e-Science in a Cloud computing environment [12]. For example, service monitoring usually occurs at the SLA monitoring stage. How can SLA monitoring collect the resource usage information to monitor associated parameters related to service level objectives? The SLA monitoring component needs to interact with a service-side / resource-side resource usage or QoS measurement component that is responsible for the acquisition of resource usage data and QoS measurement data. An important issue is what monitoring protocol should be employed? Possible monitoring protocols include the polling protocol, publish/subscribe or call-back mechanisms etc., each of which has its own advantages and disadvantages. Apart from the acquisition of monitoring information, another important issue for SLA monitoring is its functional modules. For example, the SLA monitoring should contain the following functional sub-modules: (i) Term Interpreter, (ii) Violation Evaluator and (iii) Report generator. How to develop and integrate them?

2.2.3 Service Metering

Service metering plays a fundamental role in service computing, as QoS-assured service and load balancing all require metered services to be delivered. This involves creating a generic metric model which can be used in different service instances.

In order to meter the service usage, appropriate metrics should be defined to measure the service usage. For example, in the European Edutain project, the running total of the amount of incoming / outgoing bandwidth consumption are metered to measure the hosting service for the on-demand resource provisioning purposes [14]. The SLA manager needs to be able to retrieve usage information from functional services (e.g. job services), record the usage and optionally constrain and/or bill for the usage. Different functional services are required to report usage of different measurable quantities. For example, a job service needs to report CPU usage but a data service needs to report usage of disc space. These measurable quantities, known as "metrics" needs to be represented by URIs, etc. All these aspects of SOA present research challenges and opportunities.

3 Scientific Workflow: Enabling Automatic, Efficient and Reproducible Research

Scientific workflow management systems, such as Kepler [15], Taverna [16], Triana [17], Pegasus [18], ASKALON [19], SWIFT [20] and Pipeline Pilot [21], have demonstrated their ability to help domain scientists on scientific computing problems by synthesising data, application and computing resources. Particularly, their main capabilities include: 1) Usability: scientific workflows support easy execution process expression and sharing among users; 2) Automation: diverse tasks and resources can be integrated into one scientific workflow and executed without user interaction; 3) Efficiency: workflow engines can automatically enable efficient execution based on certain targets and optimisation rules; 4) Reproducibility: workflow executions could be fully or partially reproduced, so that scientists can better understand the computation process.

Scientific workflows have been an important component since the beginning of e-Science, which has resulted in many successful applications in various disciplines. A few applications are listed here for illustration purposes. Taverna is used in the caGrid platform⁷, which aims to share data and integrate services of many clinical or research fields including cancer research, public health and neuroscience. Using Kepler to build and execute metagenomic analysis processes, CAMERA⁸ provides the microbial ecology research community with a comprehensive data and bioinformatics tool repository. Pegasus is employed as a workflow engine in the Southern California Earthquake Center (SCEC) CyberShake project⁹, an analysis activity to compute probabilistic seismic hazard curves around the southern California area. Galaxy¹⁰ is a Web-based platform for accessible, reproducible, and transparent computational biomedical research, on which users can easily run tools and workflows.

3.1 Workflow Research Aspects and Community Practice

Scientific workflows can be used in various settings including local, Web or distributed applications and focus on different targets, such as reproducibility, execution efficiency, generality and sharing. Thus research on scientific workflows also varies in different projects. More complete descriptions of scientific workflows can be found at [22], [23], [24], and [25]. We only try to discuss some key and common research aspects here.

Figure 3: Conceptual view of the enzyme design process [26].

⁷ <http://www.cagrid.org>

⁸ <http://camera.calit2.net/>

⁹ <http://scec.usc.edu/scecpedia/CyberShake>

¹⁰ <http://galaxy.psu.edu/>

A real world e-Science application in Computational Chemistry is illustrated in Figure 3. This application takes quantum mechanical theozymes as inputs, takes three discrete steps, and is validated by experiments in the final stage. Although the conceptual process looks simple, there are many challenges when trying to standardise it through workflow automation and eliminate unnecessary human interactions. Firstly, the corresponding executable workflow should be easy to build by correctly modelling the data, tools and the logic of the process. Secondly, the composed workflow should run efficiently and adapt to different execution environments since users might have various data sizes and different computation resources. Thirdly, execution provenance information might need to be recorded for future query and analysis because one workflow often needs to be executed many times with different parameters and re-executions could be very expensive. These challenges are more difficult when the problem scales up. Even for the three-step enzyme design case, one whole computation required for all 226 reference datasets (called the scaffold) includes thousands of tool invocations, takes hundreds to thousands of CPU hours, and results in about seven million enzyme designs.

Figure 4: Lifecycle of scientific workflow and its research aspects.

By summarising and generalising the above challenges, we think the lifecycle of the scientific workflow consists of three basic phases (Figure 4): composition, execution, and learn. Workflow composition is about how to build workflow applications based on user requirements. It usually needs to follow the specification defined by the chosen workflow system. Different workflow systems support different workflow models and usability. After composition, workflows will be executed to get their results. In many cases, workflow composition only defines the conceptual logic in a workflow, not concrete computation resource information. So composed workflows will be scheduled to actual computational resources and the adaptability of a workflow system is defined by its capability to execute the same workflow on different computation resources. After workflow execution, users can obtain workflow execution results and provenance for each data product, and share their workflows with others. As a special type of knowledge, a whole workflow or part of a workflow could be shared in future workflow composition. As feedback, the knowledge learnt here will help users to better understand their scientific problems and improve corresponding workflows. For instance, users can determine execution bottlenecks from provenance information and update their workflows to be more efficient.

3.1.1 Workflow model

Workflows commonly include three types of components: tasks, control dependencies and data dependencies [26]. A workflow needs to follow certain dependency logic for its execution. The dependency logic is typically expressed via *control flow*, *data flow*, or a *hybrid* of both. For control flows, also known as control-driven workflows, explicit control structures (such as sequence, loop, condition and parallel) are employed. In data flows, or data-driven workflows, data dependencies are used to describe the relationships among tasks. Two tasks are only connected when the downstream task is to consume data from outputs of the upstream one. The hybrid method adopts both control and data dependencies for powerful and easy logic description.

Besides the above common components, the current main research aspects of workflow models focus on their support for some specific semantics: implicit/explicit parallelisation, data streaming, continuous time, discrete event, higher-order functions (e.g., iteration, map, reduce), and so on. Each semantic expresses a special problem with certain characteristics, and normally results in a special reusable building block or template in the workflow model. For instance, originating from functional programming [29], map is a higher-order function that applies a given function to each element of a list and the function execution for the elements could run in parallel. By having a map template in workflow that accepts lists or arrays as input, users can build a sub-workflow for the given function. When the workflow executes, the sub-workflow will automatically run against each element of its input list/array.

Many scientific workflow systems, e.g., Kepler, Triana and Taverna, support hybrid dependency modelling. Some workflow systems are specialised to a few special semantics for their project requirements. For example, workflows in the LEAD project focus on discrete events and real time responsiveness to enable dynamically adaptive weather analysis and forecasting [27]. Inheriting from Ptolemy¹¹, Kepler can support many different semantics including data flows, discrete events, process networks, and continuous time. Recently, more and more workflow systems [30], [31] are supporting Map-Reduce so that large datasets can be efficiently processed through distributed data parallel computation.

3.1.2 Usability

Generally, the usability of a workflow is important in all phases of its lifecycle. Here, we only focus on the ease of use during workflow composition. Many workflow users are domain scientists whose main interests are their scientific problems rather than programming. So it is very important for them to easily build workflows based on their requirements.

Since many common tasks are used in many workflows, they are usually pre-defined as components to facilitate workflow composition. The tasks could be either within one research domain, such as sequence alignment for bioinformatics, or across domains, such as data transfer. Domain knowledge is critical here to build and organise a repository with reusable and configurable components.

Besides the organisation of reusable components, task organisation is also important within a workflow, since workflows could easily get too complicated. Hierarchy is a good way to organise a workflow. If a workflow has multiple levels, a lot of details can be embedded in lower levels, leaving the upper levels much clearer. A component built from other components and processes is often called a *sub-workflow* or *composite task*. Users do not need to know what is inside a sub-workflow if they are not interested. Further, sub-workflows can also be shared amongst many workflows.

¹¹ <http://ptolemy.eecs.berkeley.edu/ptolemyII/>

To simplify workflow composition, many workflow systems including Kepler, Taverna and Triana, provide graphical user interfaces and support drag-and-drop to add tasks into workflows. They also support hierarchies within a workflow via sub-workflows. With more and more available task components and pre-built workflows, they need to be managed effectively. Many semantic techniques including ontology and tagging are used in task and workflow repositories to help their searching and organisation.

3.1.3 Scheduling

To realise automation, workflows need to be executed on computational resources to obtain results. Workflow scheduling maps the tasks and data in a workflow to real computational resources. An important research aspect of workflow scheduling is determining an appropriate scheduling solution in order to meet expected targets or constraints, in which many factors are to be considered. The first one is the workflow model. The workflow scheduler should check each task in the workflow to determine how, when and where to run it, and follow the overall logic of the workflow to find possible efficient ways to allocate pending tasks to resources. The second factor is the information about data to be processed and tools to be executed. To execute each task, the relevant data and tools have to be accessible by the computational resource. There will be an efficiency challenge when a large amount of data or tools need to be moved across the network. The third factor is the computational resource information: the capability of each resource determines whether it can execute a task and how fast the execution will be; the number of available computing resources is important when there are many pending tasks that can run in parallel. The dynamism of available resources may force scheduler to re-schedule workflow when the resource availability changes, which is quite common for Grid computing environments. The fourth factor for workflow scheduling is the expected targets or constraints, which are decisive in choosing scheduling algorithms. Targets or constraints could be the user perspective, e.g., the minimal execution time and acceptable execution time deadline, or the resource perspective such as maximal resource utilisation.

Generally, workflow scheduling in distributed environments is an NP-hard problem and there are no optimal solutions in polynomial time [28]. A lot of near-optimal workflow scheduling algorithms have been proposed for diverse requirements and execution environments [35], [36], [37],[38]. They are mainly classified as approximation algorithms and heuristic algorithms [35]. Approximation algorithms [39] are used to find approximate solutions to the optimisation of problems and have provable bounds for the objectives. Heuristic algorithms [40] are able to produce an acceptable solution to a problem in many practical scenarios, but for which there is no formal proof of its bounds.

Recently, new targets or constraints have been proposed and studied to meet the new developments of scientific computing problems or available resources. To deal with growing big data challenges, some workflow scheduling algorithms have been proposed to consider data placement and movement optimisation [32], [33]. Along with the growing popularity of Cloud computing, many new scheduling algorithm studies have taken into account monetary cost targets or constraints [58], [59].

3.1.4 Adaptability

Many scientific workflow systems have certain adaptability features since the composition and execution phases of a workflow can be separated and even operated by different users. During the execution phase, the information about data and computation resources is obtained either from user settings or automatic detection. Then workflow execution engines will use this information to find an appropriate way to interact with the resources.

Workflow adaptability mainly deals with the following types of problems: 1) Resource Allocation: one workflow task could be executed by either resources from multiple accessible

candidates, e.g., computer clusters or Web services; 2) Protocol Selection: one data/tool in a workflow may be accessed/invoked through multiple possible protocols, such as different file protocols and remote procedure call protocols; 3) Computation Unfolding: one task or sub-workflow could actually need to be iterated or parallelised during execution based on different conditions, e.g., file number of a directory or parameter value. Adaptability differs on how easily a workflow can be adapted to different resources and the range of adaptable resources.

As a layer between users and resources, workflow can provide certain abstraction to deal with different underlying resources. To enable adaptation, some workflow systems including Pegasus and Swift support abstract workflows. These abstract workflows only describe minimal task and data information, and the logical dependencies between them. During the execution phase, they will be mapped to executable workflows based on available computation and data information.

3.1.5 Sharing

Reflecting knowledge by describing computational processes as workflows for a certain scientific problem, workflows are valuable and often shared with other researchers in the same domain. Research areas within workflow sharing include workflow versioning, organisation, authorisation, validation, searching, interoperation, and so on.

To facilitate workflow sharing, Web 2.0 and Semantic Web techniques have been employed. As a good representative, MyExperiment [41] is a popular Web site for sharing workflows. It can be used to search, share, tag and review workflows, and to communicate with other users, such as creating and joining groups, finding people, and building your own reputation. Besides sharing complete workflows, component sharing within one workflow system is also important so that users can contribute their work and utilise others' work. Both Kepler and Galaxy support component sharing and importing through Kepler actor repository¹² and Galaxy tool shed¹³, respectively.

Sharing workflows across different workflow systems brings challenges in workflow interoperation, since each workflow system originally only supports workflows written following its own specification. A common solution is to invoke workflows written in other specifications via a standard protocol, such as Web services. This is a coarse-grained solution since workflows with other specifications are encapsulated in black boxes. To realise white box or fine-grained interoperation, a shared or standard workflow model is needed so that workflows with different specifications can be converted to the shared workflow model before their execution. Several efforts have been underway to facilitate workflow sharing between workflow systems [34]. A recent project, called SHIWA¹⁴, tries to support both coarse-grained [42] and fine-grained workflow interoperability [43]. It provides a workflow repository and a simulation platform to enable sharing and execution of workflows built from several workflow systems.

3.1.6 Provenance

Provenance is metadata for “the derivation history of a data product, starting from its original sources” [64], which plays a critical role in scientific workflows by recording necessary details (e.g., data inputs/outputs, parameter values, execution times and locations) throughout their execution. Provenance can help scientists in many ways: 1) reproduce scientific workflow executions since reproducibility is a key requirement for scientific research; 2) avoid full or partial workflow re-execution especially for some workflow

¹² <http://library.kepler-project.org/>

¹³ <http://toolshed.g2.bx.psu.edu/>

¹⁴ <http://www.shiwa-workflow.eu/>

executions that are expensive or time-consuming; 3) facilitate new findings by querying and comparing provenance information of multiple executions; 4) identify bottlenecks of existing workflows for possible optimisation. Research areas within workflow related provenance include the provenance data model, provenance data query, provenance recording and querying performance, and distributed provenance.

The standardisation of provenance has made a lot of progress recently. The Open Provenance Model (OPM)¹⁵ allows provenance information to be exchanged between systems and shared amongst users. A W3C Provenance Working Group¹⁶ has also been created as part of the Semantic Web Activity (see Section 4.2.2), which is to “support the widespread publication and use of provenance information of Web documents, data, and resources”.

Many other research aspects of provenance remain as active research topics, especially on how to wisely and efficiently save provenance for data-intensive applications[129],[130], and dealing with new provenance challenges in Cloud environments[131],[132].

3.2 Workflow Open Issues and Opportunities in e-Science

Along with the evolution of related techniques and higher expectations from users, there are always open issues for scientific workflows. Many research questions mentioned in the last sub-section are still open. Here, we will emphasise on four emerging and important issues for scientific workflows.

3.2.1 Workflow Scheduling in the Cloud

Because of its abundance and scalability characteristics, Cloud is becoming a popular environment for scientific workflow execution and brings new challenges for workflow scheduling. In Cloud environments, especially for commercial Cloud resources, e.g., Amazon EC2¹⁷ and Microsoft Azure¹⁸, there are usually usage charges for running applications, taking into account the allocated processor time. When using these resources, users naturally want their usage to be both execution and cost efficient. Since execution and cost efficiency conflict with each other at most times [35], practical user requirements are commonly expressed as *objectives with quality of service (QoS) constraints*, such as 1) the objective of minimising the total completion time with a budget constraint (namely the monetary cost limitation for the total execution) or 2) the objective of minimising the budget with a total completion time constraint. More detailed information of QoS can be seen in Section 2.2.1.

The challenge here is to find proper workflow scheduling algorithms on Cloud resources that can help to automatically find optimal or near-optimal resource usage plans that meet user objectives and QoS constraints for their applications. Most existing workflow scheduling solutions are only able to minimise the total completion time [35], [36], and [37]. They have to be extended or modified to support budget as a constraint or objective. Due to the control or data dependencies, a task in a workflow cannot start before all its dependent tasks are finished. It is quite common that not all designated processors are busy during workflow application execution. One processor in the Cloud could be reserved in different ways, such as from the workflow start time to the workflow finish time, from the time it start to execute a task to the time it finishes all tasks designated to it, or only the time when there are tasks

¹⁵ <http://openprovenance.org/>

¹⁶ <http://www.w3.org/2011/prov/>

¹⁷ <http://aws.amazon.com/ec2>

¹⁸ <http://www.windowsazure.com>

running on it. These ways will result in different budget calculation formulae, and have different difficulties to implement them.

Since more and more scientific workflow applications are running in the Cloud, QoS constraint based workflow scheduling in Cloud environments is becoming an active research area [58], [59]. Based on different application and Cloud characteristics, more specific workflow scheduling approaches in Cloud environments are expected to be developed.

3.2.2 Big Data Processing in Cloud

Traditionally, scientific workflows move data to local or remote locations where scientific applications employed in the workflows are deployed, before the execution of the applications. If one workflow contains applications deployed on different sites, there have to be multiple data movements during the workflow execution.

The above solutions will meet challenges when dealing with data deluge situations that are common in many scientific domains [45]. DNA sequence data in biology is a representative example. With the introduction of the next-generation sequencers, e.g., the 454 Sequencer, there has been a huge increase in the amount of DNA sequence data. For example, the Illumina HiSeq 2000 can produce 600 billion base pairs per run [65]. When the data sizes are large, times for data movements will be significant and result in very inefficient workflow executions.

We discuss how big data or data intensive workflow applications could be dealt with based on the recent advances in programming models and Cloud infrastructure. At the programming model level, many distributed data parallel patterns identified recently provide opportunities to facilitate big data applications/workflows [44]. Typical patterns and their supporting frameworks include MapReduce [46], All-Pairs [47], Sector/Sphere [48], Hadoop¹⁹, and Stratosphere²⁰. The advantages of these patterns include: (i) a higher-level programming model to easily parallelise user programs, (ii) they follow the “moving computation to data” principle instead of traditional moving data to computation, which can reduce data movement overheads, (iii) support for data distribution and parallel data processing on multiple nodes/cores, (iv) good scalability and performance acceleration when executing on distributed compute nodes, (v) support for run-time features such as fault tolerance and security, (vi) simplification of the difficulty of parallel programming in comparison to traditional parallel programming interfaces such as MPI [49] and OpenMP [50].

At the infrastructure level, Cloud computing and storage provide opportunities to facilitate data-intensive applications and workflows. Their advantages include: 1) on-demand resource provision for scalable applications, and 2) dynamic binding of Cloud storage to virtual Cloud computing resources. In Infrastructure as a Service (IaaS) Clouds such as Amazon EC2, we can create a customised image containing deployed applications and application-specific data storage in the Cloud beforehand. Then, virtual instances can be instantiated based on the image, and attached to application specific data storage. In this way, applications and data can be coupled during execution without additional data transfer.

Cloud infrastructure also supports the above data parallel patterns very well. Therefore, with Cloud and data-parallel programming models, we will have good workflow execution efficiency by not only having applications close to data, but also supporting good application scalability.

¹⁹ <http://hadoop.apache.org>

²⁰ <http://www.stratosphere.eu/>

Figure 5: A framework for data intensive workflow applications in the Cloud.

Figure 5 illustrates a framework for data intensive workflow applications using the MapReduce distributed data-parallel pattern in the Cloud. Both Master and Slave processes of the MapReduce pattern are running on virtual instances in the Cloud. By using a distributed data-parallel engine, such as Hadoop, we can achieve data partition and distribution on the Slave instances. Also the workflow/application to process the data can be locally accessible through customised Slave images. With the support from the distributed data-parallel engine, the workflow on the Master instance can manage the overall execution on the Slave instances. This framework can also be adjusted to use other distributed data-parallel patterns mentioned above.

There have been some scientific applications that run in the Cloud and utilise data-parallel patterns to achieve good scalability [51], [52], [53], [128]. There are also some general research [54], [55], [56], [57], and specific systems such as Oozie²¹, Azkaban²² and Cascading²³ supporting data-parallel patterns in workflow. We believe there will be more and more data-parallel workflow applications to deal with the data deluge in e-Science.

3.2.3 Workflow Execution in Hybrid Environment

From a user perspective, scientists usually want to focus on their domain specific problems rather than the evolving distributed computing techniques. Many workflow users are unfamiliar with existing distributed computing frameworks, e.g., Globus²⁴ and Hadoop, and requiring knowledge of these types of systems effectively prevents adoption of distributed computing.

From the computational resource perspective, it is reasonable that different distributed environments will co-exist for a long time and work in a hybrid way [60]. There have been various distributed computing techniques, such as Cluster computing, Grid computing and Cloud computing, that can be utilised to accelerate workflow execution. Each distributed computing technique has its applicable context. A workflow might be executed across multiple types of distributed environments.

A big challenge here is how to make the distributed execution of scientific workflows adaptive to a hybrid distributed environment. The theoretical enzyme design workflow in Figure 3 and [26] is a good representative instance for this challenge. Computational methodology for enzyme design has been developed using quantum mechanics and molecular

²¹ <http://yahoo.github.com/oozie/index.html>

²² <http://sna-projects.com/azkaban/>

²³ <http://www.cascading.org/>

²⁴ <http://www.globus.org/toolkit/>

dynamics. One whole computation required for all input data could contain millions of atomic jobs and take months to execute on one single CPU core. Furthermore, the same workflow needs to be executed many times with different parameters. The available computing resources include a few local or remote clusters, a Grid environment and commercial Cloud resources like EC2. No additional fees need to be paid for each usage of the cluster and Grid resources since they are owned by users or shared between them. Yet it is still an appealing idea to use commercial Cloud resources as additional resources for some jobs with some monetary cost when the in-house or shared resources are busy and the execution of the whole workflow is expected to finish soon.

By providing an abstraction layer between user application and underlying distributed environments, workflow systems are good candidates to facilitate scientific application execution in hybrid environments. Workflow systems will need to interact with different underlying system and scheduling tasks in a workflow to be executed on proper resources. When running workflows in hybrid environments, a good workflow system needs be able to 1) adapt itself to reuse existing workflows and minimise user involvement, 2) find the best scheduling solution based on user targets and/or constraints, 3) minimise the overhead or difficulties of using hybrid environments, such as additional data transfer and access control management.

There has been some initial work dealing with application execution in hybrid environments of Cloud and other types of resources where workflow is one key component [60], [53],[61] . It is expected that more capable workflow systems for hybrid environments will appear, and more scientific applications will utilise them to achieve adaptation and efficiency.

3.2.4 Delegation in Workflow

When composing multiple Cloud services into workflows, problems are commonly faced when these services require access control. Consider a situation in which a scientist needs to process a large amount of data, which are held on a remote server that requires the scientist to authenticate in order to be granted access. The data are to be processed on a processing resource in a different location from the data server. To avoid having to download large amounts of data to their own machine, the scientist wishes the processing resource to access the data directly on their behalf. The workflow is created using a website.

The Figure 6 illustrates this situation. There are three servers involved: the Web server, the processing resource and the data server. Each requires the user to authenticate, but the user only interacts directly with the Web server. The user's authority must be *delegated* down the workflow chain, ultimately permitting the processing resource to act on behalf of the user to access the data.

Figure 6 A typical scientific workflow involving remote data processing. The services employ access control (indicated by padlocks), requiring the user to *delegate* their authority to services in the chain to act on their behalf. Large data transfers to and from the user's machine are avoided.

The MashMyData project²⁵ investigated this problem for a particular instance of this kind of scientific workflow. The data server and processing resource were part of pre-existing infrastructure in the Centre for Environmental Data Archival in the UK. They are accessed through protocols that are widely used in the geospatial community: OPeNDAP²⁶ for data access and the Web Processing Service (WPS)²⁷ interface for the processing resource. Both service instances required short-lived EECs (End Entity Certificates) to grant user access. The Web portal used OpenID for user authentication. The multi-step delegation problem was solved in the project using a combination of credential translation and the well-established proxy certificate mechanism from the Grid Security Infrastructure²⁸.

OAuth²⁹ is an alternative technology for delegation that has gained traction particularly in the commercial Web environment. Well known sites such as Twitter make use of OAuth to enable third parties to act on the user's behalf to access resources belonging to them hosted at another site. Where proxy certificates adopt an approach of delegation by impersonation, OAuth has the ability to enforce more fine grained control delegating a limited authorisation. Work in the CILogon [62] project has shown how OAuth can be used to protect a short-lived credential service and avoid the need for custom SSL middleware required by consumers to correctly verify proxy certificate delegation chains.

Another increasingly-common situation in which delegation is required lies in the use of “brokering” services. Brokers are used to provide new interfaces atop existing services, or to aggregate and filter results from many services [63]. For this brokering approach to work with services that require authentication, a solution to this “delegation problem” is required.

Both the OAuth and certificate-based solutions require thorough testing in a production environment. Additionally, there are key challenges remaining in ensuring that delegation systems strike an appropriate balance between security and user-friendliness. It will be very important to ensure interoperability among different access control mechanisms in order to allow the user to combine and process data seamlessly.

²⁵ MashMyData <http://www.mashmydata.org>

²⁶ OPeNDAP, <http://www.opendap.org>

²⁷ Web Processing Service, <http://www.opengeospatial.org/standards/wps>

²⁸ At the time of writing this paper, a paper for addressing this multi-step delegation problem is in preparation.

²⁹ OAuth, <http://oauth.net/>

4 Semantic Web in e-Science

The increasingly large volumes of scientific data along with the research models underpinning the data need to be accessible, verifiable and re-usable across the world to provide evidence based solutions to the grand challenges of health, natural disaster prediction and social policy making. For the technologists involved in the value chains of those solutions, the same e-Science data and research models could facilitate generation of significant economic benefits.

The Semantic Web provides technologies to allow information on the Web, which is human-usable, to be described with richer semantics to enable them to be computer processable. Once the information is computer usable it allows it to be linked together to facilitate more effective discovery and accessibility of Web-based resources. As a result, there has been great interest in applying the Semantic Web techniques to facilitate sharing and discovery of scientific knowledge, particularly in the Life Sciences [66]. In principle, the Semantic Web technologies could allow e-Science data to be linked together so that the underlying models can be applied to wider sets of data, and specialist local models can, in turn, be linked together to provide evidence based solutions to the broader and more important problems noted before.

4.1 Overview of Semantic Web Techniques and Concepts

An overview of the key Semantic Web techniques and concepts relevant to e-Science is presented in this section.

4.1.1 Semantic Web Languages

The notion of “Semantic Web” is underpinned by a number of languages defined and/or standardised by the World Wide Web Consortium (W3C)³⁰ to enable formal description and efficient querying and analysis of concepts, terms, and relationships within a given domain of knowledge. These Semantic Web languages include Resource Description Framework (RDF)³¹, Turtle³², N-Triple³³, the Web Ontology Language (OWL)³⁴, RDF Schema³⁵ and SPARQL³⁶. In general, the intended purposes of these languages can be categorised as **Resource Description**, **Knowledge Representation** and **Querying**. At present, RDF is commonly used for describing Web-based resources identified with Uniform Resource Identifiers (URIs), while OWL is used to represent semantically rich knowledge and information models (typically in the form of an ontology) underpinning such Web-based resources. RDF Schema can also be used for defining information models and data structures but with relatively less semantic expressivity. SPARQL enables querying resources encoded in any of the Semantic Web languages for resource description.

³⁰ W3C - <http://www.w3.org/>

³¹ RDF - <http://www.w3.org/RDF/>

³² Turtle – Terse RDF Triple Language - <http://www.w3.org/TeamSubmission/turtle/>

³³ RDF Test Cases (N-Triples) - <http://www.w3.org/TR/rdf-testcases/#ntriples>

³⁴ OWL 2 Web Ontology Language - <http://www.w3.org/TR/owl2-overview/>

³⁵ RDF Schema - <http://www.w3.org/TR/rdf-schema/>

³⁶ SPARQL Query Language for RDF - <http://www.w3.org/TR/rdf-sparql-query/>

4.1.2 Linked Data

Figure 7: Linked data principles: client-dependant resource identification through HTTP URI and resource retrieval through content negotiation. (Source: <http://www.w3.org/TR/cooluris/>)

The two core functionalities of today's World Wide Web are the ability to identify and link documents using the HTTP protocol. These two elements are simple to implement, widely deployed, and have ubiquitous client support. As a result, they provide an obvious model for moving beyond text and documents to a Web of data – where related datasets as well as documents could be linked together and exposed through the HTTP protocol as Web-accessible resources. The 'linked data principles' [67] adopt this model by using URIs to identify data objects (or the real-world 'things' that they represent), and creating a Web of data by linking together related data objects. While HTML provides the *lingua franca* for the Web of documents, RDF plays that role for data (Figure. 7). Common to both is the use of HTTP to access information; linked data also recommends a human-readable representation e.g., HTML, if accessed via a Web browser, using 'content negotiation'³⁷ (Figure. 8). The adoption of the four elements of linked data – i.e. URIs, RDF, HTTP, links between data – has already led to a massive 'linked data Cloud'³⁸ connecting hundreds of datasets and billions of individual data items [68]. Figure 8 below illustrates how HTTP URIs could be used to uniquely identify and provide information about a real-world concept "River Thames" as a Web-based resource that could also be linked to other related Web-based resources (also identified with unique HTTP URIs) through the RDF vocabulary, "rdfs:seeAlso".

³⁷ A web server returns a representation of a resource based on the HTTP-Accept header of a client request.

³⁸ <http://linkeddata.org>

Figure 8: Use of HTTP URIs and RDF vocabulary to identify, expose and link related Web-based resources.

4.2 The Role of Semantic Web in e-Science

We outline some of key roles that Semantic Web can play in facilitating sharing and enrichment of e-Science data and the underlying research models through effective data publishing and integration, accurate and comprehensive data provenance capture as well as efficient data curation and annotation.

4.2.1 Data Publishing and Integration

Traditionally, the formal scientific output in most fields of natural science has been limited to peer-reviewed academic journal publications. Datasets have been and continue to be archived, but the scientific focus remains on the final output, with less attention paid to the underlying research model or workflow containing the chain of intermediate data results and their associated metadata, including provenance. This has effectively constrained the representation and verification of the data provenance to the confines of the related publications, as well as limiting the scope for cross-disciplinary integration of data and models to facilitate better science and analysis [69].

The notion of publishing data sets has become widely discussed within the data management community. By adopting a mechanism whereby data is “published”, that is made available to the public as an item of record. When undertaken within a strong publication process, data publication has some advantages:

- Delivers a definitive reference copy of the data with some guarantee of its stability over time, available to the experimental scientists themselves and the wider scientific community.
- Allows the data to be accessible (with suitable access conditions) for validation and reuse.
- Allows a notion of quality assurance to be applied to the data, by for example third party review.
- Potentially, it allows the data collector to gain credit for the collection of the data.

As a further driver, external organizations such as publishers (e.g. International Union of Crystallography³⁹) and funders (e.g. the UK Engineering and Physical Science Research Council⁴⁰) are requiring the deposit of data in a reliable archive to maintain the evidential basis for research finding and to make the data available as a reusable asset.

As noted above, the field of data publishing is still relatively underexplored despite gaining momentum of late due to the drivers and motivations cited above. Notable endeavors in data publishing include some of the NERC and JISC funded projects, such as the OJIMS and CLADDIER projects, as well as the more recent NERC SIS⁴¹ data citation and publication project which builds on the former two projects. The approaches outlined in the OJIMS [76], [77] and CLADDIER [75] projects are very general. CLADDIER investigated differing methods for publishing datasets, and discussed the requirements for the peer-review of data, as well as proposing a structure for human readable citation strings. OJIMS took the case of an overlay journal for data publication and created a demonstrator journal, investigating the business case for operating it on a long-term basis, as well as surveying the proposed user community about their opinions on data publication and their use of data repositories [69].

The Semantic Web techniques, particularly linked data has the potential to extend the work done in both of these previous projects, and take it down to a more detailed level, focusing more on the complete trail of provenance associated with a dataset; as mentioned before, provenance information is essential for data validation amongst other things. This was illustrated in the recently completed JISC-funded ACRID project⁴². In essence, ACRID defined RDF ontologies to describe the data, metadata and workflows associated with complex climate science datasets, and publish them using a combination of Digital Object Identifier (DOI)⁴³ and linked data compliant data re-use standards (e.g. OAI-ORE⁴⁴) to enable a seamless link between a publication and the detailed workflow associated with the corresponding datasets. Besides data publishing, there are also well understood examples where the integration of data and models in different disciplines provide compelling benefits. For example, in bioscience the linking of genome data to protein data and epidemiological health data could help deduce potential drug treatments. In environmental sciences, the global linking of seismic and other geological data supporting global predictive models of earthquakes, volcanic activity and tsunamis provide access to more accurate models than local models. In structural sciences, the linking of experimental data with the derived molecular structures to the uses of the compounds provides insight into novel applications. Beyond linking data, there are clear examples where models can be integrated - for instance, in biology where models of cell operation can be linked into organ models and then in turn into complete system models to provide an evidence-based chain of prediction up through the layers of description.

The Semantic Web provides technologies to define the terminology used to describe data and models (ontology languages such as OWL), along with languages to link data and models ("linked data" technique underpinned by representation languages such as RDF). In addition, it facilitates the development and provision of the inference engines and data management tools to use those languages to reason over scientifically significant features of data and models such as provenance and quality.

³⁹ <http://www.iucr.org/>

⁴⁰ <http://www.epsrc.ac.uk/Pages/default.aspx>

⁴¹ NERC SIS: <http://www.nerc.ac.uk/research/sites/data/sis.asp>

⁴² Advanced Climate Research Infrastructure for Data (ACRID) - <http://www.cru.uea.ac.uk/cru/projects/acrid/>

⁴³ The Digital Object Identifier (DOI) System - <http://www.doi.org/>

⁴⁴ Open Archives Initiative Object Reuse and Exchange (OAI-ORE) - <http://www.openarchives.org/ore/>

4.2.2 Data Provenance

We discuss the data provenance from the Semantic Web perspective. Data provenance will be discussed further in Section 6.3.7. As discussed above, the publication of data makes available the evidential basis of scientific claims within paper, allowing them to be validated for accuracy and reused within new contexts. However, providing access to “raw” data alone is of limited value. Data need to be provided with additional annotation of its context - e.g. how it was collected, the conditions of its environment and processes applied to refine the data- in order to allow the proper interpretation of data. Such data provenance information can be considered as annotations within controlled vocabularies and links to other related resources involved in the process - e.g. related data, instruments, software, people - distributed within the community. Semantic Web and linked data principles are thus ideal to provide representation for such networks of related information. This has been explored in such projects as ACRID [69] which published provenance data as linked data using the OAI-ORE aggregation format⁴⁵. As mentioned previously, the W3C has established a Provenance working group⁴⁶ to provide a common representation within the Semantic Web for provenance information, and Ontologies such as the SPAR ontologies are being developed to represent the relationships between research results⁴⁷. However, linked data while providing a basis for publishing provenance, needs to be augmented with mechanisms to handle aggregation of different resources in a secure and sustainable manner[70], [69]. Such aggregations, known as Research Objects⁴⁸ could be considered the units of publication beyond traditional journal articles within a linked data environment.

4.2.3 Data Curation & Annotation

The ability to publish e-Science data and underlying research models as well as the associated provenance is reliant on the effective curation of the associated metadata. These metadata include information about the data context and meaning, integrity of the preservation process and important assumptions about the target user community that may change over time. The underlying curation process would need to involve capturing of accurate metadata at crucial junctures of the data life-cycle, quality assurance, efficient management (e.g. versioning) of the metadata captured, and finally storing it, ideally in a medium that is suitable for efficient querying and dissemination of the metadata. Without effective curation the metadata may become out of step with the data, which may lead to inaccurate and/or incomplete provenance description of the data [73].

As noted above, the Semantic Web languages, particularly the knowledge representation language OWL, provide a means of developing suitable ontologies to define and capture accurate metadata about data. However, effective data curation is more than just capturing and publishing metadata; it also involves adding value to data – e.g. through the means of annotation.

Annotation in the digital world has long been recognised as an effective means of adding value to digital information. It can in effect, help establish collaborative links between data providers and data users. However, annotation without the intended context may become meaningless. For example, an annotation may be used to label particular components of a scientific workflow with descriptive text, which may contain values of some attributes

⁴⁵ Open Archives Initiative Object Reuse and Exchange <http://www.openarchives.org/ore/>

⁴⁶ W3C Provenance Working Group http://www.w3.org/2011/prov/wiki/Main_Page accessed 18 Dec. 2011

⁴⁷ Semantic Publishing and Referencing Ontologies (SPAR) <http://purl.org/spar/> page accessed 18 Dec. 2011

⁴⁸ <http://www.researchobject.org/>

associated with those components. These attribute values alone, without the correct association with the corresponding context, would be meaningless. For complex and dynamic environmental datasets, it may be useful for users to be able to annotate specific features or attributes for collaborative analysis or interpretation, for instance in an emergency response scenario.

Linked data and other related Semantic Web technologies have the potential to represent annotations as Web-based resources linked to their corresponding scientific datasets and/or workflows, while providing accurate identification of the context referenced and detailed description of the underlying knowledge. This potential for Semantic Web based approaches to annotation was considered in e-Science even before the recent impetus in the adoption of the linked data and Semantic Web related technologies which have recently been extended with vocabularies for provenance⁴⁹ and annotation⁵⁰. For example, in 2004, the Conceptual Open Hypermedia Services Environment (COHSE)⁵¹ was used by the myGrid⁵² project to semantically annotate provenance logs generated by the various bio-scientific experiments considered by the project. These provenance records were also conceptually linked together as a hypertext Web of provenance logs and experiment resources, based on the associated conceptual metadata and reasoning over these metadata [73].

In addition to annotation, long term preservation also needs stable references to resources so that guarantees can be made on the persistence of those resources over time [72], and initiatives such as PURL⁵³ and DOI⁵⁴ to provide dependable identifiers for resources within the Web.

4.3 Key issues and challenges

As noted above, Semantic Web technologies, in particular the linked data technique, provide the opportunity to integrate research data and models to develop evidence-based solutions to various social, political and environmental problems. However, these solutions require individual researchers to believe in the benefits of that integration, to understand the technologies which can be used, and to trust that those technologies will be assimilated sufficiently widely by the community they identify with to justify them putting effort into adopting them.

There are also several other caveats to effectively sharing linked resources using URI and RDF. The chief amongst these is the necessity of a specific community data model, or ‘RDF vocabulary’ or ontology. While RDF provides the base representation for linked data, this is not enough to specify the internal structure of any specific dataset (much as HTML provides a flexible structure for a huge variety of Web page content). As noted by Tim Berners-Lee [67], “Different communities have specific preferences on the vocabularies they prefer to use for publishing data on the Web. The Web of Data is therefore open to arbitrary vocabularies being used in parallel. Despite this general openness, it is considered good practice to reuse terms from well-known RDF vocabularies...” Unfortunately the most well-known RDF vocabularies/OWL Ontologies are too generic to describe domain specific knowledge or models, such as climate science related information model. Most of these RDF

⁴⁹ PROV-O: The PROV Ontology <http://www.w3.org/TR/prov-o/>

⁵⁰ Open Annotation Collaboration <http://www.openannotation.org/>

⁵¹ COHSE - <http://cohse.cs.manchester.ac.uk/>

⁵² myGrid project - <http://www.mygrid.org.uk>

⁵³ Persistent Uniform Resource Locators <http://purl.oclc.org/docs/index.html>

⁵⁴ Digital Object Identifier <http://www.doi.org/>

vocabularies/OWL ontologies are concerned with social networking (FOAF⁵⁵), blogs/wikis (SIOC⁵⁶), thesauri (SKOS⁵⁷), software projects (DOAP⁵⁸), etc. Communities need to come together and agree on the Ontologies that are appropriate to their own domain so that they can share concepts. Again, the bio-science community are leading the way, with efforts such as the Gene Ontology⁵⁹ which is seeking to agree on consistent description of gene products. These can then be used to annotate resources and allow searching and exchange of data.

Further, the ability to link resources may not necessarily translate into the ability to effectively exchange and share those resources, unless the linking and exchange formats are either the same or equally common within the associated community. RDF, the recommended linked data format, though gaining increased adoption, is not a commonly used format for exchanging data within every scientific community in the world. For example, the geospatial community predominantly relies on the Geography Markup Language (GML)⁶⁰ representations of the ISO 19100 series models along with other geographical data formats, such as NetCDF for encoding and exchanging environmental data rather than the Semantic Web resource description languages, such as RDF. So, in communities and domains within which Semantic Web techniques have yet to garner major uptake, the linked data approaches to describing and publishing data would need to support commonly used data exchange formats (e.g. GML for the geospatial domain) in addition to RDF.

In addition, a linked data service should integrate with existing data sources without needing to make substantial changes to the underlying infrastructure. For example, it may not be desirable to significantly modify an existing Web Server serving up external data from a third party database; or to replace it with a linked data service to provide linked data representations of these data. What might be more efficient and practical in this scenario is to implement a linked data service that wraps the Web Service and leverages it as a “proxy” data source for exposing linked data.

4.4 Scaling the Semantic Web

e-Science problems are increasingly becoming those of “Big Data”, as data acquisition and storage increases, there are two aspects of scalability which need to be tackled: those of volume data (and associated metadata); and the heterogeneity of different data semantics – interdisciplinary science in particular needs to address issues of combining diverse data sources. The Semantic Web can be fully beneficial to e-Science, by overcoming the technical problems of defining common vocabularies, publishing data on-line, marking up on-line data with those vocabularies, linking the data and models, but also the social problems of getting scientific communities, who do not identify themselves with each other, to work together towards a common, highly abstract goal. However, it is not as yet clear that the current generation of Semantic Web tools scales up to the big- data challenge.

Cloud computing technology has the potential to solve some of these technical problems of the Semantic Web to allow it to scale in proportion to big-data. For instance, harmonising and interconnecting large heterogeneous datasets through a common ontology could lead to the need for a great deal of computing power, which could be difficult to provide through traditional centralised computing platform. The notion of Cloud computing is

⁵⁵ The Friend of a Friend (FOAF) vocabulary - <http://xmlns.com/foaf/spec/>

⁵⁶ Semantically-Interlinked Online Communities (SIOC) - <http://sioc-project.org/ontology>

⁵⁷ Simple Knowledge Organization System Reference (SKOS) - <http://www.w3.org/TR/swbp-skos-core-spec>
⁵⁸ <http://code.google.com/p/baetle/wiki/DoapOntology>

⁵⁹ The Gene Ontology Project <http://www.geneontology.org/>

⁶⁰ Geography Markup Language <http://www.opengeospatial.org/standards/gml>

intended to enable traditionally controlled software applications, whether a standalone desktop application or a Web service, to evolve into “on-demand” remotely accessible Web applications. In principle, this should provide a suitable platform for developing simple but powerful solutions for creating and publishing complex linked data in the Cloud. The potential for offering or exploiting “Linked Data as a Service” in the Cloud is increasingly being recognised. For example, Haase et al. [74] presents a technological platform, namely Information Workbench, which supports self-service linked data application development. In general, this platform aims to support discovery and exploration of linked data resources. More pertinently, the underlying architecture of the platform adopts a “Data-as-a-Service” paradigm in order to facilitate virtual integration and processing of the linked data resources. In principle, this provides the potential for deploying the platform based on Cloud technologies.

As highlighted in [78], using Cloud computing platforms and technologies in conjunction with Semantic Web technology could be mutually beneficial. On the one hand, use of Semantic Web tools and the underlying metadata models and ontologies as part of increasingly popular Cloud services is likely to broaden understanding, and thus, adoption of the Semantic Web tools and ontologies – currently a barrier to wider adoption of the Semantic Web paradigm [see Section 4.3]. On the other hand, Semantic Web-conformant metadata models have the potential to enable semantically richer and standardised service description, more effective analysis, and, as a result, optimum utilisation of Cloud computing services. According to [78], this could effectively enrich the “semantics of the Cloud computing landscape itself”. However, this is still an emerging field, so the mutually beneficial possibilities of integrating Semantic Web with Cloud computing need to be explored in greater detail.

5 From Grid computing to Cloud: e-Science perspective

The Cloud computing paradigm, although it has emerged as a current IT trend, is not a completely new concept. It has an intrinsic connection to the well-established Grid Computing paradigm and other relevant technologies, such as utility computing, distributed systems as well as social computing. This section aims to discuss the evolution from Grid computing to Cloud computing by focusing on the comparison of essential characteristics of both. A vision on how e-Science could evolve in the Cloud computing is presented.

5.1 e-Science Evolution: from Grid to Cloud

Grid Computing [79] was proposed as an infrastructure to enable the sharing of computational resources for e-Science. The development of Grid technology has been focused on building Grid middleware as a platform such as Globus [80], gLite [81], UNICORE [82] to dynamically organise geographically distributed heterogeneous resources across multiple organisations to form a uniform computational utility. Although Grid computing focuses on the challenge of organising distributed resources to offer a high performance computation platform, resources are normally provided and administrated at project or organisation level. It pays little attention to the effective use of computation platforms shared by a large community with many divergent applications. It is expected that Infrastructure as a Service (IaaS) of Cloud computing can meet this goal.

Fundamentally, Cloud computing and Grid computing have the common goal to reduce the cost of computing, increase reliability and flexibility by transforming computers / platform, software, and applications from something that we build, buy and operate to services operated by third parties. Differences between Cloud and Grid can be distinct from the following aspects.

Business Models: A clear difference is that Cloud computing was developed with a clear business driver. The core technologies of Cloud Computing, such as elastic resource management and multi-tenancy, have been focused on the efficient provision of a centrally managed vast system in an on-demand manner based on a pay-as-you-go pricing principle. The pay-as-you-go model is indeed the core concept which has driven the Cloud computing model to become a phenomenon sweeping across the Internet today. On the Grid computing side, although technologies have been there for more than a decade to organise distributed resources into a virtual data centre, there has been no clear business model in providing such an aggregated resource pool nor mechanism in supporting its provision to the general public, as Grid computing was mainly adopted by research organisations and government labs to solve “grand challenge” problems. Although some research projects such as GridEcon [85], ArguGrid [86] did propose some concepts that try to enable economy-awareness in Grids, they did not generate impact on a sufficiently large scale to drive the Grid computing towards the direction of true utility computing.

Computing Models: Grid and Cloud have clear differences in computing models. The design rationale of Grid computing focuses on resource aggregation which leads to a job centric computing model. Most Grids adopt batch-scheduled work model, which relies on a local resource manager that manages the pre-registered computing resources. User submitted jobs are queued, scheduled and executed over the managed resources. When a job is scheduled for execution, the resource allocated is fully committed to the executing job. In Cloud, when providing computing resources, the computing model is resource centric and based on the well-known time-sharing concept. Rather than occupying computing resources through job execution, users virtually own the computing resources for the time that is required. The computing resources are provided as virtual machines on-demand and directly managed by the end users. This is what we call the computing elasticity. This principle can also apply to other Cloud service provision. For example, the database can be virtualized to enable the sharing of a physical database management system (DBMS) by a large user community.

Programming Model: The programming model for Grid systems inherits from parallel and distributed computing. It focuses on the scalability to leverage large amount of resources. Distributed workflow management systems are widely used in the Grid computing environment as programming tools since they enable a flexible application-level programming paradigm which better suits the major Grid users: the scientists. Unlike in the Grid environment where resources are usually provided for free, Cloud services provision is highly economic driven. This requires programming in Cloud to consider balances between the quality of the results and their costs. Cloud applications normally have resource management, metering and billing functions built-in to achieve better quality/cost trade-offs, as discussed in section 2.2. Some of the Cloud services are even non-deterministic; non-deterministic services refer to those applications that deliver multiple results with different qualities with respect to different cost constraints [83]. Developing a new algorithm framework and a programming paradigm for such a “pay-as-you-go” computation is an exciting research field.

To summarise, a Cloud offers a shared computational environment on the Internet with the following five features: (1) on-demand self-service, (2) broad network access, (3) resource pooling, (4) rapid elasticity, and (5) measured service. Some of the features are shared with, or precisely originated from, Grid computing, such as broad network access and resource pooling. It is the emphasis of efficient provision of shared resources to a broad community for a wide range of applications that made Cloud computing move from the Grid view of “aggregated resources as a platform for solving a big problem” to a true realisation of the utility computing where “aggregated resources as an utility shared by a large community for solving many problems”.

5.2 e-Science 2.0: e-Science in Cloud

It is a quite restricted view if we only consider Cloud as a new computing resource provision model where hardware resources of a data centre are provided as services based on a pay-as-you-go pricing model (i.e. IaaS). The essence of Cloud computing goes far beyond IaaS. With the vision of providing all the resources on Internet as services, the Cloud computing model forms an industrial structure and a business paradigm for Internet-based information society. Recent Internet computing technologies, such as those sometimes collectively described as Web 2.0, contribute to the key components of the Cloud computing eco-system where new paradigms such as social computing and collective intelligence are built up to realise a new era of Internet computing. Such a Cloud computing eco-system provides a new underlying infrastructure for e-Science.

Table 1 Evolution of e-Science Underlying Infrastructure

Web 1.0 (Machine as the platform)	Web 2.0 (Internet as the platform)
Personal Computer	Mobile Devices (IPad, iPhone)
Enterprise Computing	Cloud Computing
Relational Database	RDF/DataSpace
DB Application	Force.com (SaaS)
IDE (Eclipse)	Google AppEngine
Installer	AppStore
Britannica Online	Wikipedia
Personal Web Sites	Blogging
Address Book	Facebook
Content Management Systems	Wikis
FTP	Dropbox
Media Player	Youtube

Table 1 illustrates the evolution of the underlying infrastructure for e-Science from a local to an Internet-based platform. Such an evolution has a profound impact on the development of e-Science. If we view the achievement of e-Science in the past 10 years as changing the way scientific research is undertaken, the future development of e-Science will bring changes into the organisation mechanism, social computing and knowledge dissemination model for Internet-based scientific research.

Hence we envisage that “e-Science 2.0” will provide a platform to enable new scientific research practice, where the scientist can share raw experimental results, workflows, nascent theories, claims of discovery, and draft papers by forming a dynamic community on the Web to collectively conduct the research (i.e. open science). The research results can also be openly shared to the general public (as services), and knowledge can be further evolved during its open application.

Figure 9 illustrates the layered framework of “e-Science 2.0” in the Cloud. The e-Science infrastructure service is built upon the existing Grid infrastructures to take the full advantages of them. It turns the underlying Grids into Internet based services which can then be

consumed, exchanged and even traded. Through adopting such a layer, many isolated (both conceptually and geographically) Grid sites would be combined into one or several scientific resource Clouds which are accessed and utilised by a global user community. On top of the e-Science infrastructure, a scientific service layer provides domain specific services such as scientific workflow management services, genome analysis tools and etc. which would facilitate scientific research processes. In the Grid setup, many of these tools and services have been designed and implemented and some of them can be quickly transplanted into the Cloud environment directly as the technologies used are not affected by the paradigm change.

Figure 9 Layered Framework for e-Science 2.0 Services

As the core of e-Science 2.0, collaboration and community services provide a virtual collaboration environment over the internet for e-Science 2.0 users. Giving support from the underlying layers, e-Science users are not only able to share scientific data such as experiment results and publications but also able to work on the same experiments / research topics using collective intelligence. Also, people can form virtual organisations and communities. Those communities can use the services provided at this layer to spark new modes of discovery, innovation, learning, and engagement that will accelerate the transformation of science. It helps to produce innovative theory, modelling, and simulation that are tightly linked to experimental research and to education.

The framework shows that by using a conceptually centralised Cloud science platform, scientists, researchers and even general citizens are able to undertake experiments, derive results, and share knowledge in a collaborative manner. Within such a context, scientific activities can have more direct and greater influence on the world. Three main goals would be achieved through the use of the Science-as-a-Service in the Cloud: (1) Transparency in experimental methodology, observation, and collection of data; (2) Public availability and reusability of scientific data; (3) Public accessibility and transparency of scientific communication. A prototype Cloud, namely, IC Cloud [84] has been developed following the proposed architecture to demonstrate this collaborative scientific research environment of e-Science 2.0.

5.3 Open issues and challenges of e-Science in Cloud

In this sub-section, we discuss the two technical issues of e-Science in Cloud. Actually, the most important and critical issues for enabling e-Science for open research are social rather than technical (e.g. open data), but this is beyond the scope of this paper.

5.3.1 Cloud interoperability

From a technical perspective, issues such as interoperability of different Cloud providers, porting existing e-Science services onto the Cloud, supporting new scientific service models for the Cloud presents challenges to e-Science in Cloud. The essence of these problems is that each vendor's environment supports one or more operating systems and databases. Each Cloud contains its own hypervisors, processes, security, data storage model, networking model, Cloud API, and licensing models, etc. Though some of these issues involve specific business models, technology advances perhaps can help to address these challenges and these also present research opportunities.

5.3.2 Virtualisation

As a key enabling technique of Cloud computing, virtualisation provides a simulated execution environment on top of physical hardware and systems. Using current virtualisation techniques, users normally can work on the virtual machines as if they work on physical machines.

Virtualisation can have a positive impact on e-Science. First, virtual machines provide idle test beds for e-Science. When a virtual machine is started, it can be administered and configured as needed without affecting the physical machine. If the virtual machine stops working for some reason, the user can simply delete the virtual machine and start a new one. Secondly, it is easy to get on-demand execution environments via virtualisation without investment in new hardware. Different scientific applications often have differing software requirements (e.g. operation system type, operation system version, tool version, library version, etc.) and hardware requirements (e.g. CPU, memory, disk, network). With virtualisation, virtual machines with various configurations can be created on the same physical machine. Further, customised virtual machines can be shared between researchers to save time in system installation and configuration.

e-Science could also be affected by some disadvantages of virtualization. For example, the performance of scientific applications might be compromised because of the overhead of having the additional virtual machine layer. However, experiments show the overhead is very small [87], [123]. The performance of scientific applications might fluctuate and be unpredictable over time. It occurs when multiple VMs are running concurrently on one physical machine (also called multi-tenancy), causing the workload of one VM to affect the performances of others. We expect this problem will be alleviated with the improvement of isolation techniques of virtualisation.

It is easy to set up a virtual Cloud computing lab for a research group, community or university. We can either directly use public Cloud environments, such as Amazon EC2 and Microsoft Azure, or use a private Cloud environment by installing Cloud computing software, e.g. OpenStack, Nimbus and Eucalyptus, on existing local machines or clusters. There are two main reasons for choosing one way over another. The first is whether the group wants to pay upfront cost for buying hardware and maintain it. The second is the utilisation of the Cloud computing resources. Public Cloud is more economically suitable for occasional usage and private Cloud is more suitable for constant usage. More detailed experience of setting up a

virtual Cloud computing lab can be found in [88], and cost and performance comparison for the scientific application in Cloud can be found in [89].

No matter which way we use to set up a Cloud computing lab, its components are basically the same. We need hardware for virtual machines to be hosted, and proprietary or open source Cloud software to build and manage virtual machines. Users of the Cloud computing lab need to build, access, monitor and use the virtual machines. They can do it through a Web portal, command line tool, or Web service.

In general, challenges for implementing associated e-Infrastructures using Cloud computing can be classified into two main categories: data and computation. The next two sections (section 6 and section 7) are to discuss them in details.

6 Cloud-based Storage and Preservation in e-Science

Cloud storage is commonly defined as a virtualised on demand storage service [90]. Such services offer highly scalable storage accessible to both individual researchers and research institutions based on a pay-per-use model. Cloud storage providers can exploit the economies of scale of large data centres built with high density, efficient energy usage and low staffing requirements, to provide efficiencies that are difficult to achieve at an organisational level, for all but the largest organisations. Commodity Cloud storage has the potential to dramatically reduce the requirement for research institutions to maintain internal storage facilities and data centres for all but the most sensitive or high-performance applications (Examples where more specialised facilities might be required are medical applications involving real patient data, or particle physics, which has a requirement for storing and processing huge volumes of data.).

There is an increasing trend towards data-intensive research in many fields. Such data can be generated, for example, by instruments, sensor networks or simulations, which require ever increasing storage capacity. Traditional areas of e-Science such as particle physics are already well-supported by infrastructures such as Grid computing, which are well-adapted to dealing with the huge volumes of data. Cloud storage opens up research opportunities for data intensive research in areas such as social sciences, humanities, and environmental science, which have not traditionally been supported by custom e-research infrastructure such as Grid and do not have the scale or funding to support such facilities. Indeed the availability of Cloud computing is opening up new research fields amenable to quantitative methods that were previously dominated by qualitative approaches.

The rate of growth in storage technology, at 25% per annum, is being outstripped by the rate of growth in data at 60% per annum, highlighting the requirement for both short-term as well as longer term storage solutions [93]. Placing data in the Cloud provides greater opportunities for collaboration through sharing of working and archived datasets. This should be combined with mechanisms for data cleansing and preservation and providing suitable identifiers and provenance information.

Despite the great potential in using Cloud resources to support e-Science, there are considerable technical, legal, governance and trust issues to be overcome. User research with academics across a number of disciplines carried out by the Jisc-funded Kindura project [91] has highlighted that many users are unwilling to share valuable research data with commercial Cloud providers. Issues such as low bandwidth for uploading large datasets, fear of loss of service, potential breach of confidentiality agreements for sensitive or personal data, and lack of understanding of the costs involved are recurring themes. We explore some of these issues in the following sections.

Many of the issues associated with Cloud usage are not fundamentally issues with the technology itself, but rather concern more subjective issues such as trust, as well as legal,

governance and financial matters. However, further technological development can assist in resolving such problems.

6.1 Cloud storage in e-Research

Cloud storage supported by a standard set of Cloud-based computational tools can potentially provide support for the full research data lifecycle from data capture to archival of research results. A particular incentive for placing data in the Cloud is ease of sharing across institutions and national borders, and the potential for discipline-based repositories and computational tools running in the Cloud. The potential role of Cloud storage and supporting tools, in the research lifecycle is illustrated in Table 2.

In particular, Cloud provides a long term solution for archiving source data and experimental results. Since Cloud storage is accessible over the Web, it provides a platform for data sharing and reuse, and for supporting the open access agenda. Currently, academic papers typically contain only summarised results and graphs describing research datasets. The availability of on-demand storage means that source data and final results as well as intermediate datasets can be retained in online repositories. Placing data in Cloud storage also makes possible to make use of shared curation services that enable longer term preservation.

Lifecycle stage	Description	Supporting tools
Data capture	Harvesting of data from external sources and repositories, capture from instruments and sensors, human entry (e.g. citizen cyber science ⁶¹), capture from mobile devices.	<ul style="list-style-type: none"> • Data cleaning, • Quality control, • Annotation, • De-duplication, • Provenance capture.
Analysis	Software used to perform analysis of the source data, producing intermediate and final datasets.	<ul style="list-style-type: none"> • Visualisation, • Metadata creation logging and annotation, • Data backup.
Sharing and collaboration	Support for sharing the data in the Cloud, across distributed projects or research networks.	<ul style="list-style-type: none"> • Access management, • Usage monitoring, provenance tools
Preservation and archival	Support for data curation and quality control, data repositories, creation of preservation metadata, retention policies.	<ul style="list-style-type: none"> • Cloud-based repositories, • Curation tools, • Identifiers to support data citation.

Table 2 Cloud storage to support the research data lifecycle

Many Cloud storage providers offer an associated compute services that enable free internal transfers of data. Repeatability and verification of experimental results is often difficult to achieve without access to all relevant data and suitable computational resources.

⁶¹ <http://www.citizencyberscience.net/>

6.2 Grid versus Cloud-based Storage

Grid systems such as iRODS [94] and its forerunner SRB (Storage Resource Broker) [97] are gaining acceptance in traditional e-Science fields such as particle physics, providing a mechanism for distributed storage across multiple research facilities. The Grid storage model is based on storage quotas in contrast to the pay-per-use and on-demand Cloud model. The barriers to entry to using Grid resources are higher, requiring a steep learning curve and are typically the domain of highly IT-literate users, as well as considerable capital investment. Hence, with the exception of specific research domains, Grid has not had the impact on research that might have initially been expected [104]. Grid technologies have also not found wide acceptance in the commercial sector.

Grid storage systems have been typically been employed in conjunction with high-performance computing (HPC) infrastructures, which often entail the generation of huge files that need to be shared under strict data usage policies. Systems such as iRODS enable the interconnection of distributed storage resources to enable increases in demand to be met, and for these systems to be managed as a single infrastructure. iRODS supports rules implemented as microservices to support the flow of information. Skilled users or administrators can then author or amend services to implement the desired functionality.

In an IaaS form, Cloud offers a relatively simple Web interface that is accessible to a wider range of less technically oriented users. Cloud infrastructure services are typically available through CRUD (*Create, Read, Update, Delete*) operations implemented as REST-based services over HTTP. In order to preserve the simplicity of the interface, metadata are used to configure the services themselves. Payment for services by credit card widens the accessibility to individual researchers.

Cloud storage systems such as Amazon S3 have a property called eventual consistency [108]. This describes when changes committed to a system are visible to all participants. Eventual consistency allows a time lag between the point the data is committed to storage and the point where it is visible to all others, which theoretically has no upper bound. Hence, in practical terms, a delay must be employed between writing and reading the same data object, which may be a disadvantage compared to Grid systems for some applications.

Commercial Cloud computing offers little transparency for institutions or users in understanding how their data is stored and managed beyond the limited information provided by service providers. Cloud computing in particular offers an attractive option to “small science⁶²”, that is communities of researchers who require to share information and resources, providing high scalability and availability. Cloud can also provide a key enabler for research students to access computing resources on an ad hoc that an academic institution would be unable to resource through internal infrastructure.

A further issue that inhibits the use of Cloud storage in the academic sector is network bandwidth. Although research organisations are connected by high speed networks, these networks are currently not well connected to the networks used by commercial Cloud providers. Applications such as DNA sequencing, requiring massive parallel computation, are increasingly being run in the Cloud [96]. Moving large datasets over low bandwidth connections is often impractical, making migration between Cloud providers difficult. Often data transfers of large datasets are carried out by physical shipping of portable storage devices.

At an institutional level, the rapid growth in use of Cloud technology in enterprises has led to a widening skills gap in the areas of virtualisation and Cloud computing [95].

⁶² <http://onlinelibrary.wiley.com/doi/10.1029/EO085i027p00260-03/abstract>

6.3 Challenges and Opportunities in Cloud-based Storage

Challenges and opportunities of Cloud-based storage are identified as follows.

6.3.1 Deployment and Integration with Cloud Storage

Many organisations in the commercial sector are already making use of Cloud storage and compute services in their IT infrastructure [95]. Hybrid Cloud solutions are gaining in popularity as a way of combining multiple forms of storage including internal storage, private and public Clouds. This can be used to hide the physical storage location from end users, provide appropriate governance structures to ensure the integrity and security of sensitive and business-critical data. A disadvantage of this approach is the additional middleware that is required to mediate between in-house systems and external Cloud infrastructures, in terms of additional cost, reliability and sustainability.

6.3.2 Ease of use

Fundamental to the uptake and exploitation of the full potential of Cloud storage for e-Science is the provision of easy-to-use tools. Many researchers are already implicitly making use of Cloud storage via applications such as Google Drive, iCloud and DropBox. Cloud storage is also a popular medium for backup of research data as well as personal data.

A widening of data intensive research to areas such as social sciences and humanities is resulting in a greater use of Cloud storage by less IT proficient users. These research communities make heavy use of standard desktop applications such as spreadsheets and databases. Microsoft Research has developed Cloud-based Excel services that connect with Azure to enable spreadsheet calculations to be performed on large datasets [98]. The Jisc-funded Biophysical Repositories in the Lab (BRIL) project⁶³ demonstrated how a mapped folder on a user's PC can be used to capture, annotate and preserve data from experimental workflows across multiple applications. Linked data objects are created automatically that can be deposited in Cloud repositories with minimal user input.

Overall there is a requirement for greater integration of Cloud storage into existing research tools and applications, to hide the raw Web service interfaces and the management of storage from the user.

6.3.3 Cloud Storage Marketplace

Cloud IaaS storage providers such as Amazon S3, Microsoft Azure and Rackspace provide services that are generic and highly scalable but at the expense of being general purpose and not targeted to the requirements of the academic research community. Indeed it would be difficult for more niche providers to match the economies of scale achieved by global providers. Organisations such as EduServe and the National Grid Service (NGS) in the UK have explored the market for resources dedicated to the academic sector, at the expense of limited scalability and a quota-based allocation model^{64, 65}. For institutions and end users, understanding and comparing the available Cloud storage services and choosing the most appropriate provider for a given task is complex. The market is evolving rapidly with new

⁶³ JISC Biophysical Repositories in the Lab project (BRIL), <http://www.jisc.ac.uk/whatwedo/programmes/inf11/digpres/bril>.

⁶⁴ Eduserve Managed Hosting and Cloud, <http://www.eduserve.org.uk/hosting>

⁶⁵ UK National Grid Service, <http://www.ngs.ac.uk>

players entering and provider offerings changing frequently. Planning tools are required to enable better understanding of likely costs of hosting services in the Cloud so that the costs are more transparent. An example of such a tool is the Cloud Adoption Toolkit [92]. In order to make decisions about migrating data storage into the Cloud or migration of data between Cloud providers, tools are also required to predict the costs, in order that realistic comparisons can be made.

6.3.4 Standards for Cloud Storage

The development of Cloud technology has been substantially driven through the commercial sector, with the interfaces and protocols by major providers such as Amazon being regarded as de facto standards. Lack of formally agreed standards raises the risk of vendor lock-in making it difficult for users to move their data to the most cost effective provider. This includes differing metadata formats for expressing storage configuration and permissions. The cost of bandwidth for data transfers and the time taken to migrate data are also important factors.

The standardisation agenda around Cloud storage has been driven by a number of standards bodies, as well as by the European Commission, through its ICT research funding programmes and the Siena Initiative⁶⁶. Standardisation and adoption of Cloud services and APIs is still at an early stage. The Storage Networking Industry Association (SNIA) launched the Cloud Storage Initiative (CDMI)⁶⁷ in April 2009. Its aim is to understand requirements for managing data in the Cloud and the possible technical solutions. CDMI defines the functional interface that applications use to create, retrieve, update and delete data elements from the Cloud. CDMI provides a REST API for management of data objects including files and database tables, as well as metadata schema to configure storage characteristics such as replication and retention. A standard notion of container is provided with associated metadata. Accounts allow management of user permissions and third-party sharing. CDMI Capabilities enable characterisation of the available storage operations.

The Open Cloud Computing Interface (OCCI) comprises a set of open community-lead specifications delivered through the Open Grid Forum (OGF)⁶⁸. OCCI is a Protocol and API for various management tasks associated with Cloud. For a more comprehensive list of organizations involved in standards activities related to Cloud computing, see [105]. Further effort is required in standardisation, particular to achieve buy-in and adoption by the larger commercial providers.

The IEEE has provided support for research activity in Cloud and Grid computing for many years. In 2010, the IEEE Cloud Computing Initiative (CCI) was launched to provide a more coordinated response to the opportunities in this area. The IEEE initiative is focused on three main areas: Cloud infrastructure, platform and application interfaces, units of measurement and registration authorities. IEEE 2301⁶⁹ examines interfaces (application, portability, management, interoperability), file formats and operations conventions, and groups these into multiple logical profiles, to address the requirements of ecosystem participants (Cloud vendors, service providers and users). IEEE 2302⁷⁰ addresses interoperation and federation of Clouds, and is aiming to develop specifications in such areas as protocols, directory services, trust and namespace authorities. The standard aims to create a

⁶⁶ <http://www.sienainitiative.eu>

⁶⁷ CDMI standard, SNIA, <http://www.snia.org/cdm>

⁶⁸ <http://www.ogf.org/>

⁶⁹ <http://grouper.ieee.org/groups/2301/>

⁷⁰ <http://grouper.ieee.org/groups/2302/>

marketplace for Cloud services that is more transparent to users. A number of universities are developing real-world test beds to support the development and validation of the standard.

One approach to resolving the incompatibility between Cloud APIs is provided by the DuraCloud open source middleware developed by DuraSpace⁷¹. This enables connection to multiple Cloud storage providers through a common REST API. DuraCloud provides a plug-in framework for connecting Cloud storage services. Each Cloud provider still requires metadata in custom formats for configuration of storage settings and permissions. Further work is also required to integrate billing systems from multiple providers into this framework.

Standards have emerged that aim to provide interoperability between Grid and Cloud systems. The Open Grid Forum (OGF) GLUE32 [106] standard provides an information model for describing grid and Cloud entities. The DMTF Common Information Model⁷² (CIM) provides a common definition of management information for systems, networks, applications and services, and allows for vendor extensions, enabling the exchange of semantically rich management information.

6.3.5 Costs and Cloud Capability Models

Public Cloud providers use different and continuously changing models for charging for resources. As the market for Cloud storage becomes more competitive, the complexity of such pricing schemes is likely to increase, mirroring the situation with other utility services such as power distribution. Users typically pay both for storage usage at a Gb/hour rate, as well as for inward and outward data transfers at possibly differing rates.

For current users, there is a lack of transparency in the costing information, making comparison between provider offerings difficult, also in relation to comparing prices with storage performance. Thus there is a requirement for effective cost models and storage provider brokerage. This in turn needs comparable cost and capability models. An XML schema for characterising storage provider capabilities and costs has been proposed by Ruiz-Alvarez and Humphrey [99]. Existing schemas already exist that could be adapted or extended for this purpose including the Resource Specification Language (RSL)⁷³ and the Job Submission Description Language (JSDL). The SoaML⁷⁴ (Service oriented architecture Modelling Language) specification and the CloudML project⁷⁵ whose goal is to develop extensions to SoaML for the deployment of applications and services on Cloud for portability, interoperability and reuse, provide an alternative approach for describing storage resources.

Individual users paying for public Cloud storage using personal or institutional credit cards raise financial concerns for institutions and funders wishing to control their expenditure and demonstrate value for money from taxpayer support. It is difficult for institutions to effectively account for outgoings on Cloud resources. Potential economies of scale cannot be exploited by gaining cost reductions through pooling of resources. Cloud providers use differing non-standard accounting mechanisms, which hinders integration with institutional finance systems.

The current charging models for commercial Clouds do not readily support sharing of data on Cloud sites by enabling end users of data to pay for the costs of data they are downloading

⁷¹ DuraCloud, DuraSpace, <http://duracloud.org>.

⁷² www.dmtf.org/standards/cim

⁷³ <http://www.globus.org/toolkit/docs/5.0/5.0.0/execution/gram5/pi/>

⁷⁴ <http://www.omg.org/spec/SoaML/>

⁷⁵ <http://aws.amazon.com/cloudformation/>

in a simple and transparent manner. Thus the development of more flexible and standardised billing and security management technologies is required.

Requests for funding of Cloud storage resources may become increasingly common on grant proposals, rather than contributions to capital spending on hardware. Again there are potential economies of scale that can be exploited. UK funding councils such as EPSRC and NERC are increasingly mandating retention of research data for periods of 10 years or more [100], [101] with a similar trend in other countries. Hence there is a requirement for low-cost long term storage options within the Cloud ecosystem. Strategies are required for migrating content from higher cost low-latency storage to cheaper long-term storage are required to reduce longer term storage costs. Automated or semi-automated management of retention policies is required to ensure that content that is no longer required can be deleted.

Despite its relatively high cost compared to internal storage, the Cloud provides an attractive option to research institutions as it enables spending on storage to be moved from capital to operational budgets. This reduces the risk in making large one-off investments in infrastructure. Tracking of Cloud storage usage across an institution is critical to understand usage and expenditure patterns. The use of Cloud also challenges the way in which central IT departments are funded. As pay-per-use is the predominant charging model for Cloud, this could result in departments paying for IT services they consume, rather than as a fixed cost.

Overall there is a requirement for more transparent costing to enable effective comparison of providers and provider brokerage, as well as payment and contractual models better suited to the needs of e-Science.

6.3.6 Legal and Regulatory Issues

Cloud data centres of large corporate providers are distributed across national boundaries and legal jurisdictions. There is therefore currently little transparency for the user in understanding where their data is stored and who might have access to it. Service Level Agreements (SLAs) are typically highly weighted to suit the service provider. End users paying for Cloud services by credit card often do not have the necessary legal expertise to interpret the SLA agreement into which they are entering. Techniques such as Cloud bursting [118] are often permitted by such agreements, allowing storage providers to outsource data for storage to third parties. Machine-readable Service Level Agreements (SLA) in standard formats would greatly assist the user in selecting services and enable the user to be assisted by the development of automated tools. The SLA@SOA⁷⁶ project considered the development of standards for specification and negotiation of SLAs taking into account monitoring, accounting and enforcement. More about SLA can be found in Section 2.2.

The Edward Snowden revelations in 2013⁷⁷ have led to demands from businesses and citizens for increased transparency about the locations and legal jurisdictions in which their data is stored. In light of this, many governments are reviewing their data protection and privacy laws.

A parallel trend is the establishment of *vertical Cloud*, which is the optimisation of Cloud computing and Cloud services for a particular vertical, such as a specific industry sector or application domain. The Cloud provider thus offers specialised functions and options that best meet the needs of the particular industry use. Examples of such verticals are health and finance, which are subject to strict regulatory requirements imposed in different jurisdictions. This might include providing much greater control of the locations in which data

⁷⁶ <http://sla-at-soi.eu>

⁷⁷ <http://www.theguardian.com/world/the-nsa-files>

is stored and matching the legal compliance requirements of users with those of the storage locations.

Health and pharmaceuticals have been some of the slowest sectors to adopt Cloud computing due in part to security concerns in managing sensitive personal data. There has been a great deal of work on electronic health records and centralised systems that can integrate data from multiple sources, to support more efficient information management and provide rapid access to critical information in medical emergencies. Cloud provides an ideal platform for such services. Providing auditing of Cloud systems to ensure compliance with regulatory requirements is still an open research question. For example [126] proposes methods for providing third party audits of Cloud systems using cryptographic methods, without the need to transfer all the data.

6.3.7 Data quality

A strong motivation for placing data in the Cloud is for sharing of information within the research community, including source data, intermediate results and data to support published papers. Data quality is a major concern when carrying out processing of shared research datasets [114]. In particular, if the data is harvested from multiple sources, there may be inconsistencies in the representation and annotation of data fields. Cloud-based business intelligence tools, becoming widespread in the commercial sector, provide support for basic data quality operations such as data cleansing and data de-duplication. Such tools could be easily adapted to meet more effectively the needs of e-Science. For example Google Refine⁷⁸ can be run on public Clouds such as Amazon EC2, and provides a semi-automated tool for cleaning tabular data. Microsoft SQL provides a feature called a reference data service⁷⁹ that enables the validation of datasets against high quality and trusted third party stored in the Cloud.

6.3.8 Privacy, Security and Trust

Issues of privacy and trust are of particular concern to researchers, especially in fields that are processing sensitive data such as medical research. Researchers may be concerned about who might have physical access to machines such as administrators. Security breaches through malicious attacks are also a concern. In reality Cloud storage may be far more secure than locally managed hardware, due to more frequent maintenance and upgrades. If data are compromised, researchers would struggle in establishing liability against large multinational corporations. Security tracking and reporting mechanisms and standards that can be verified by recognised external bodies may assist in establishing benchmarks and compliance and building trust with users.

Whilst some e-Science disciplines allow open access to all datasets, many such as healthcare and biosciences have strict controls on the data that can be shared and with whom. Indeed this can be extended to other areas requiring the storage of personal data such as social sciences as well as commercially sensitive research. In some cases cryptographic techniques⁸⁰ can be employed to restrict access to authorised users, and also prevent unauthorised access by system administrators, when making use of public Clouds.

Trust in third party data stored in the Cloud is an important issue. For this appropriate provenance metadata is required, both to understand how and by whom the data was created and modified, as well as to understand where it has been stored and potentially corrupted. A

⁷⁸ <http://code.google.com/p/google-refine/>

⁷⁹ <http://msdn.microsoft.com/en-us/library/hh213066.aspx>

⁸⁰ A. Kumbhare, Y. Simmhan, V. Prasanna, Designing a Secure Storage Repository for Sharing Scientific Datasets using Public Clouds, <http://ceng.usc.edu/~simmhan/pubs/kumbhare-datacloud-2011.pdf>

2010 survey⁸¹ by Fujitsu Research Institute of potential Cloud customers, concluded that 88% of potential Cloud users are concerned about who has access to their data at the back-end (physical and virtual machines).

Though data provenance was discussed previously from the perspective of workflow and semantic, data provenance for trust still worth discussing. Provenance metadata is essential to determine the history of a data object. In the context of e-Science provenance [115] can be used to determine the process used to generate a dataset. This information is essential to verify the trustworthiness of the data, particularly if it is to be shared, reused or cited in a publication and to ensure the repeatability of experiments and simulations. Provenance has different but complementary interpretations in Cloud storage⁸². It can be viewed at different levels of abstraction to take into account the interaction of a data object with the layers of the stack from the application layer to the physical layer. Such provenance information can be used to record the physical and geographical locations the data has been stored, potential errors and corruption that may have occurred and accesses made to the data by human and applications. This necessitates an integrated view of provenance.

Capture of provenance in a Cloud system across multiple levels of abstraction, was addressed by the proof-of-concept Provenance Aware Storage System PASSv2 [116] system. Muniswamy-Reddy, Macko and Seltzer [117] define four principles for provenance in Cloud storage: *data-coupling* states that when a system records data and provenance, they match – the provenance accurately describes the data recorded; *multi-object causal ordering* states that ancestors described in an object's provenance exist, so there are no pointers to non-existent objects; *data-independent persistence* states that provenance must persist even after the object it describes is removed; and *efficient query* states the system supports queries on provenance across multiple objects. They also consider the implications of the violation of one or more of these options.

It is essential to be able to verify the authenticity and integrity of provenance metadata itself. In order to ensure the security of provenance metadata, it is desirable that provenance information and the data it describes have different access mechanisms. Due to the complex nature of Cloud computing infrastructures, the complexity of provenance data may be very high, resulting in higher costs for service providers and ultimately consumers. A further challenge is to make provenance metadata portable between systems of different service providers and architectures.

6.3.9 Service dependability

Service dependability and reliability are key issues for Cloud computing environments, particular for maintaining confidence of users. Due to the complex nature of the Cloud infrastructures, failures are inevitable. Failure can be caused by a number of factors such as hardware faults, software bugs, power failures, network outages, security breaches and human error. SLA contracts typically provide a minimum guaranteed availability level of service. However, the compensation provided for loss of service by the Cloud provider typically does not reflect the actual cost to a business or organisation in terms of lost revenues or damage to customer relationships. Even measurement of compliance with service quality criteria still needs to be defined in a systematic way.

⁸¹ Personal data in the cloud: A global survey of consumer attitudes. Fujitsu Research Institute: http://www.fujitsu.com/downloads/SOL/fai/reports/fujitsu_personaldata-in-the-cloud.pdf.

⁸² O. Qing Zhang, M. Kirchberg, R. K. L. Ko, B. S. Lee, How To Track Your Data: The Case for Cloud Computing Provenance, HP Laboratories HPL-2012-11 <http://www.hpl.hp.com/techreports/2012/HPL-2012-11.pdf>.

Reliability is the outcome that the Cloud-based service is fully functional from a user perspective. Constructing reliability models for Cloud infrastructures has been carried for instance in [125] and [124]. *Resilience* is the ability of a Cloud service to remain fully functional when subjected to certain kinds of failures. Resilience can be built into systems by providing flexible replacement and recovery strategies [127].

One strategy to mitigate for loss of data services is replication across multiple Cloud providers. This introduces complexities for users in managing contracts with multiple providers, as well as dealing with multiple interfaces.

6.3.10 Linked and Object Identifiers

Placing increasing amounts of data in the Cloud increases the problems of locating and identifying content items. During its lifecycle, content may be migrated across different Cloud providers and locations. This is particularly relevant for data that is stored in the Cloud for preservation purposes. In order to mitigate for loss of service or corruption, data may be replicated across multiple service providers. The identifiers and relationships between objects should be preserved even after migration. There is therefore a requirement to provide identifiers for content that are location independent and universal, that enable duplicate copies of content items to be located and maintained.

There are well known approaches for identification of digital objects such as URI [111] and URN [112]. The Digital Object Identifier (DOI) [110] system provides identifiers which are persistent and unique, has been applied in a range of publishing applications since 2000. The DOI system is implemented through a federation of registration agencies, under policies and common infrastructure provided by the International DOI Foundation which developed and controls the system. MPEG-21 Digital Item Identifier (DII)⁸³, based on URIs, provides a mechanism for identifying complex digital objects.

Linked data, based on Web technologies such as RDF and SPARQL query discussed in Section 5, provides a powerful method to link and query information stored in the Cloud [109]. Large scale systems are required to provide effective ways of searching and retrieving such semi-structured information [107].

6.4 Cloud Ecology Vision

Federation of Clouds offers the potential to combine the storage capabilities of multiple data centres and service providers to provide a high degree of scalability as well as reducing risk in reliance on a single operator. Such federation requires further progress on standardisation in such areas as APIs, service metadata, billing and SLAs in order to reduce the management complexity. Federated authentication, authorisation and identity management are prerequisites. Topologies for management of federated Clouds remain to be fully explored and might range from Clouds with a master-slave approach with strong central management functions to peer-to-peer networks with devolved management.

The European Commission has launched a number of projects in the Seventh Framework Programme to provide for interoperability of Cloud infrastructures in order to reduce vendor lock-in and enable services to run seamlessly across multiple Cloud platforms and the

⁸³ MPEG-21 standard, <http://mpeg.chiariglione.org/standards/mpeg-21/mpeg-21.htm>

exchange of data and metadata. The Reservoir [103] and Contrail⁸⁴ projects are investigating federated Cloud infrastructures as a potential solution.

In e-Science, federated Clouds could be used to support diverse requirements of the research community ranging from high speed storage to support processing of large datasets to cost-effective long term storage of research outputs. Pooling of Cloud resources across multiple data centres in academia would enable the creation of highly scalable computing and storage capabilities.

Figure 10 illustrates an architecture for a data repository based on the use of hybrid Cloud infrastructure that combines both Cloud storage and compute services with internal organisation infrastructure. This is based on work carried out in the Jisc-funded Kindura project.

The objective of such an approach is to enable preservation and storage of research data and publications through a single end-user application and interface that insulates the user from the task managing their data across multiple service providers. Such a system would provide both storage services as well as specific preservation services such as format conversion and bit-integrity checking. Service Level Agreements would be negotiated and managed at an organisational level, relieving individual researchers of this responsibility. The architecture features a service connector into which multiple storage and compute services can be plugged. Services could be added or removed to ensure the organisation benefits from the most competitive market offerings.

A data migrator handles the storage and replication of data across the various service providers, based on the use of storage rules. The storage rules take into account descriptive metadata entered by the user, metadata that is automatically extracted from the content to be uploaded and storage provider profiles that define the capabilities and costs of each storage provider. The rules can be used to implement storage policies which determine for example the number of replications of a content collection to be retained, or whether a specific content collection can be stored outside the institution. Cost optimisation is performed to determine the most cost-effective solution for storing specific content, dependent for example on the frequency of access.

A number of different migration scenarios are envisaged:

- Automated migration of data can occur when storage providers are added or removed, or as a result of a change in storage policies. In particular, this ensures the system is resilient to loss of service of a given Cloud provider.
- Semi-automated migration that involves periodic appraisal of content and migration to a different storage tier (for example migration from disc to tape).
- Manually initiated migration that occurs for instance when a user requires a copy of their data in a certain provider to enable access for computation.

Management functions handle user accounts, accounting and billing, usage tracking and system monitoring. Providing a centralised accounting system for Cloud services would enable institutions improved oversight of use of external Cloud services and remove the need for researchers to pay for services by personal credit card, by providing a common set of mediated services. A repository application with appropriate user interfaces would provide the front end to the system, providing functionalities for uploading and ingesting content, and search and retrieval.

Such an approach enables a research organisation to implement more consistent governance in the use of Cloud resources, and to provide more robust and cost-effective

⁸⁴ EU FP7 project Contrail, <http://contrail-project.eu/>

services. It reduces the time of the individual researcher to manage their storage, reduces the risk of data loss and ensures more consistent implementation of storage policies. The downside is an additional layer of middleware, which needs to be maintained.

Figure 10: Data repository based on hybrid Cloud

7 Cloud-based High Performance Computing in e-Science

In the previous sections we have considered how the Cloud has been able to provide resources and services to research communities that may previously not have had a suitable platform available to them. We are now also increasingly seeing other communities that are already well supported starting to also see the advantage of ‘Cloud’ as a paradigm or at least an on demand service.

7.1 Open issues or questions

Let us consider for example the effects of the rise of consolidation within the research sector in the provision of HPC resources. Within this section we consider HPC as being designed to support highly parallelised applications that require some form of high performance interconnect or sharing of system resources. There are a number of other activities around the suitability of Cloud services for scientific applications. These include the EC funded Helix Nebula⁸⁵, though this is concentrating almost exclusively on high throughput (trivially parallel) applications. Within the US the Department of Energy supported Magellan report⁸⁶, which addressed broader questions around the use of Cloud for science in all forms, clearly showed the benefits of using Cloud for science, particularly around the flexibility that this paradigm brings. It also makes clear, as we do elsewhere within this report, that Cloud is not a no-effort panacea, there are significant efforts that Cloud brings that are different to traditional research computing but which cannot be ignored. Most importantly though they make it clear that

⁸⁵ <http://www.helix-nebula.eu>

⁸⁶ http://science.energy.gov/~media/ascr/pdf/program-documents/docs/Magellan_Final_Report.pdf

Cloud is another weapon in the armoury, not the only one, that workloads will be suitable for different types of infrastructure and we should try to ensure that these virtual services are available alongside other traditional hardware based services,

Within the more traditional HPC community we have moved from a proliferation of separate task level resources to the institutional sized facility. This has allowed for a significant improvement in the development of hosting facilities. This in turn raises the efficiency of the facility, and with a rising efficiency comes an increase in the size of the computational service you are able to host per kW of input power to the data centre. An interesting consequence to this though has been the movement of these facilities physically away from the researchers who will use them. Through recent investments from UK funding councils (e.g. EPSRC⁸⁷) this has been taken to a further conclusion with the rise of regional facilities (e.g. e-Infrastructure South⁸⁸, Northern8 HPC⁸⁹), of a scale where large numbers of researchers from a geographically region are able to utilise the computational capacity. At this point we have to consider the definition of Cloud computing and question if general open access regional HPC facilities can be considered as a type of Cloud. This is not the first time that this has been considered. Sun Microsystems made a set of HPC facilities available in an ‘on-demand’ way well before the ‘Cloud’ paradigm was popular. In the view of the author we can consider HPC systems in this manner as a type of Platform as a Service (PaaS) system. If we move on from this and think about the different type of commodity applications that are increasingly used by communities there is another connection that can be made. Many communities have evolved with a number of applications that are considered community benchmarks that all researchers use, with examples such as Gaussian⁹⁰, DLPOLY⁹¹ and Abaqus⁹². Each of these are utilised by their communities as black box solutions where a researcher will provide the input data and configuration files but the software itself is standard. The only optimisation for the place where the software is run is through optimisation for the type of hardware system on which the software is running. Therefore, from the user point of view these are the same wherever they are installed thereby decoupling the user from the underlying system. This therefore brings the concept of Software as a Service into play where appropriate interfaces can be made available. There are already instances of this type of system available [119].

There are key questions though that must be asked about the use of Cloud within the world of HPC research. As we have shown, there are ways in which with only minor changes to user and system owner behaviour the higher levels of Cloud (e.g. SaaS) can be considered as already existing. For the lower layer, IaaS, we have to consider what are the aims of IaaS and whether in commonly available virtualisation environments that the efficiency can be raised sufficiently to make it worthwhile more generally. We have seen with the rise of both Amazon native clusters and those from reselling services such as CycleComputing⁹³ that there are clearly situations where the use of IaaS physical infrastructure to build HPC types facilities make economic and research sense. The largest impediment to these types of facilities being used more in the mainstream of research is simply the business models of many public Cloud providers.

⁸⁷ <http://www.epsrc.ac.uk>

⁸⁸ <http://www.einfrastructuresouth.ac.uk/>

⁸⁹ <http://n8hpc.org.uk/>

⁹⁰ <http://www.gaussian.com/>

⁹¹ ftp://ftp.dl.ac.uk/ccp5/DL_POLY/DL_POLY_CLASSIC/DOCUMENTS/USRMAN2.19.pdf

⁹² <http://www.3ds.com/products/simulia/portfolio/>

⁹³ <http://www.cyclecomputing.com>

An area that is not currently considered, though is normally collocated within the research computing centre hosting the HPC facilities, is how the availability of Cloud based storage could transform the use of shared and dedicated facilities within research. Consider how for example if you are able to connect large scale simulations directly through high performance networks to large data within the Cloud. Collaborators are more easily able to access these outputs, longer term data management could be easier etc., as mentioned in other sections of this paper.

7.2 Possible trends within Cloud use

Currently IaaS Cloud is viewed as the most active area within Cloud though we have seen increasingly the emergence of higher-level layers as becoming important to research. Some of the solutions that have been supported through the Microsoft Azure platform, particularly with the VENUS-C project⁹⁴ but also with native applications, have shown that user communities find this level of abstraction useful and easier for the development of complete applications. In the longer term with the appearance of feature rich open source PaaS platforms such as Cloud Foundry⁹⁵ we are also likely to see more and more applications move to this type of system. Though currently mainly supporting only Web development, more standard higher level languages (such as C, C++ and possibly Fortran) are becoming available. This may enable easier scaling though is likely to bring new challenges. This will of course be dependent on the mathematical libraries etc. being made available though this is surely dependant on the user community rather than any underlying technological impediment.

Through the previously mentioned commodity HPC applications, we are also likely to see a shift in the method of delivery for these types of applications. This will most likely be driven by the ease of management of a SaaS solution, including simplification of what in many cases are extremely complex licensing terms and conditions, and clear knowledge of the environment in which the users are operating the software by moving the operating environment ‘in-house’, overall being able to significantly improve the customer/user experience. Example platforms that could take this direction are the highly uniform computational chemistry applications, who may enter into partnerships with hardware suppliers to provide on demand SaaS services or the development/prototyping applications such as R or particularly MatLab, which is in a good position to move to the on demand pay-per-use model as a method of simplifying their extremely complex licensing models.

Another reason to combine Cloud and HPC is that the two technologies will probably co-exist in the future. The first solution is to create virtual HPC clusters on existing Cloud environments, called HPC in Cloud [120], [121]. The second one is to create Cloud environment on existing HPC systems, called Cloud in HPC [122]. In our opinion, these two trends will probably co-exist for a long time. We will discuss these two trends from IaaS perspective and we believe the situations are similar for other Cloud service models.

The first solution, namely HPC in Cloud, is mainly for users who do not want to incur high upfront costs to setup and maintain a HPC system. More and more public Cloud providers, such as Amazon EC2 and Microsoft Azure, support easy virtual cluster creation. For example, the Compute Optimized C3 instance type⁹⁶ provides not only fast CPUs but also enhanced networking, which is necessary for HPC jobs. Also there are third-party tools

⁹⁴ EU-FP7 Project VENUS-C, <http://www.venus-c.eu>

⁹⁵ Cloud Foundry, Open Source PaaS, <http://www.cloudfoundry.com>

⁹⁶ <https://aws.amazon.com/ec2/instance-types/>

such as StarCluster⁹⁷ for easy virtual cluster setup on public Cloud. Using this solution, users can quickly get on-demand HPC resources exclusively for their own use.

The second solution, namely Cloud in HPC, is for users who already have HPC systems and want to support Cloud features in HPC. Popular open source Cloud computing software, such as OpenStack, Nimbus and Eucalyptus, can create virtual instances and clusters on physical HPC systems. These software tools can handle virtualisation, virtual instance allocation, and physical resource management very well. This approach normally results in private or community Cloud environments. As explained before, users can have risk-free test beds and easy virtual environment setup and sharing through this solution. A representative example of this solution is FutureGrid⁹⁸ where Nimbus, OpenStack and Eucalyptus are deployed in its HPC clusters for Cloud usage. Often, if a computer node is deployed within such a Cloud environment, it cannot be shared for traditional HPC batch jobs. It causes the isolation between nodes for Cloud and nodes for traditional usage. If a node can be used for both Cloud usage and non-Cloud usage, it can enhance resource sharing and utilisation rate.

7.3 Lessons learned from current use

Within current exemplar scenarios of Cloud computing use for HPC applications there are a number of different lessons that have been learnt.

Firstly that a use case application is not categorically either suitable or unsuitable for use within a Cloud, within each of the categorised National Institute of Standards and Technology (NIST) service models. Within IaaS this is much more of a sliding scale depending on the hardware that is used for the Cloud system, most importantly the connection between the virtual machine manager, the internal hardware and the network interconnection. This includes the drivers that are included in any VM to any high performance network interconnection on the system. This type of specialisation also differentiates between private and community Cloud systems an institution or community may run for themselves and public Cloud that is operated at scale by groups such as Amazon etc.

An example of this type of specialisation, where resources are more targeted for a particular community, occurs in recent work by FermiLab. In comparisons between pure bare metal systems and Cloud based solutions, they produced performance figures for parallel applications that are broadly comparable to within 5 % when using the HPL benchmark [123]. The important change was to ensure the availability of the correct drivers within the virtual devices for high performance interconnection. Indeed it has been shown within a number of current evaluations of hypervisors that their performance using hardware virtualisation support in modern CPUs is nearly equivalent to bare metal. As with other the outputs of benchmarks, performance for real applications can be different. This behaviour depends on the applications themselves since the behaviour of most IaaS Cloud systems are I/O limited. For a full performance comparison we would expect to investigate a wider range of real-world user applications to provide a full benchmark.

Within the PaaS model, the individual performance can be more easily tuned to support HPC type application models. Finally, within SaaS the user is totally isolated from any underlying hardware. As this is the most true to Cloud as a business model rather than a hardware model, it is the easiest for a provider to fulfil.

⁹⁷ <http://star.mit.edu/cluster/>

⁹⁸ <https://portal.futuregrid.org/using/clouds>

Secondly, the current large problem of legacy applications and their need to be ported to new infrastructures may be solved through the use of virtual systems. These types of applications normally have a niche set of users that are unable for various reasons to upgrade to later versions of software or different comparable products. By removing the need to maintain legacy hardware systems, the support for these applications will become easier within centres that are normally limited by their support staff as much as hardware systems. A consequence is also that it is likely that apparent processor performance will be maintained comparable with the hardware systems that the applications were designed for within the virtual environment so there should be no degradation of service levels provided.

7.4 Future vision and direction for HPC in Cloud

As we move further from the idea of small HPC systems scaled at the individual application level to larger shared systems, it is inevitable that we will also move to the situation where we are able to scale the use of applications that were previously limited by the physical system size. This leads to the situation where the user is able to separate their problem from limitations imposed by the physical size of the system on which they are trying to operate and most realistically towards higher levels of abstraction that would be otherwise possible. In these situations it is inevitable that with increasing use of community applications that we could see a move towards PaaS and SaaS rather than an expansion of IaaS in the longer term.

It is clear still that the biggest question for user communities is performance, particularly where they are possibly using Cloud systems and resources where the fundamental systems are abstracted away from them to the extent of current public Cloud systems. With results of evaluations of private Cloud systems showing that the virtualisation and flexible delivery of platforms to the user can be efficient, it is likely that we may see some research computing facilities encouraging this type of use to simplify the delivery of resources to multiple different user communities. We have seen for example within the EGI⁹⁹ a requirement to move away from supporting the single application design model of high throughput computing to a more mixed model. This has led to an internal desire to push all of the current services onto a virtualised system that can then flexibly support different user communities' requirements.

Overall within public Cloud systems though, the current providers are ahead in the provision of simple easy to use systems that are targeting a very different market. As such the academic and research community and their requirement come low down on their list of priorities due to the relatively small market size. We have though seen the importance that some providers such as Microsoft give to supporting research through their participation in EC and other large scale research activities.

Within the HPC community there has been decades of development of how applications can most efficiently exploit MPI and OpenMP to gain application performance. It is likely as per the FermiCloud experience that we will firstly see a mix of the use of Cloud provisioning of systems with direct access to hardware facilities that are difficult or impossible to share. This will lead to some questioning the use of Cloud but in this instance it is all about simple application delivery and support, not whole system virtualisation. This could for example be used to share on-board GPUs within a system between those communities that require it but still enable other communities that don't to share the resources. Overall this is still a very new

⁹⁹ European Commission e-Infrastructure, European Grid Initiative, <http://www.egi.eu>

paradigm and user communities are in many ways still finding their feet as to how best their problems can be supported.

8 Conclusions

In this paper, we presented a survey of recent research advances in using Web services and SOA, workflow, Semantic Web, Grid and Cloud computing technologies in e-Science, identifying some of the research challenges and opportunities. Our survey and discussions were centred on Cloud computing. We discussed challenges and opportunities in applying Cloud computing to service computing. We presented four emerging and important issues for scientific workflows (i.e. workflow scheduling on the Cloud, data intensive workflow applications, workflow execution in hybrid environments, and delegation in workflows). It has been shown that good collaboration between computational scientists and domain experts, and thorough understanding of their scientific domains are essential to achieve successful applications. We illustrated the role of Semantic Web in data-intensive e-Science (e.g. linked data, data curation, data annotation, and data provenance), and issues and challenges when scaling the Semantic Web to the Cloud. We described the evolution from Grid computing to Cloud computing from e-Science perspective (e.g. business model, computing model, programming model). We discussed challenges and opportunities in data storage-based Cloud, and computing-based Cloud. We also presented lessons learned and future visions on using Cloud computing in e-Science. The survey discoveries in this paper will be useful to researchers and next generation e-Science infrastructure developers in understanding future trends in e-Science, or future funding and research directions that apply ICT in e-Science more effectively.

Acknowledgements

We thank the anonymous reviewers for their constructive and insightful suggestions.

References

- [1] Yang, X., Wang, L., von Laszewski, G (2009). "Recent Research Advances in e-Science", Cluster Computing special issue, 2009. Available: <http://springerlink.com/content/f058408qr771348q/>
- [2] Yang X., Wang L., et al. (2011) "Guide to e-Science: Next Generation Scientific Research and Discovery", Springer.
- [3] Hey, A.J.G. and A.E. Trefethen, *The Data Deluge: An e-Science Perspective*, in *Grid Computing - Making the Global Infrastructure a Reality*, F. Berman, G.C. Fox, and A.J.G. Hey, Editors. 2003, Wiley and Sons. p. 809-824.
- [4] John P. Sutter, Simon G. Alcock and Kawal J. S. Sawhney, "Automated in-situ optimization of bimorph mirrors at Diamond Light Source", Proc. SPIE 8139, 813906 (2011); <http://dx.doi.org/10.1117/12.892719>
- [5] Voss A., Meer E. V., Fergusson D. "Research in a Connected World", Edited book, Available, <http://www.lulu.com/product/ebook/research-in-a-connected-world/17375289>
- [6] Zhang L., Zhang J., Cai H. "Services Computing: Core Enabling Technology of the Modern Services Industry", Springer, 2007.
- [7] Yang X., Bruin R., Dove, M. et al. "A Service-oriented Framework for Running Quantum Mechanical Simulation for Material Properties over Grids", *IEEE Transactions on System, Man and Cybernetics, Part C: Applications and Reviews*, Vol. 40, No. 3, 2010.
- [8] Yang, X., Bruin, R., Dove, M. (2010) "User-centred Design Practice for Grid-enabled Simulation in e-Science", *New Generation Computing*, Volume 28, Number 2, 147-159, DOI: 10.1007/s00354-008-0082-4, Springer
- [9] Hamre, T. Sandven S. "Open service network for Marine Environmental Data", *EuroGOOS 2011*, 4-6 October, Sopot, Poland.
- [10] Browdy S.F. "GEOSS Common Infrastructure: Internal Structure and Standards", *GeoViQua First Workshop*, Barcelona, 18, February, 2011.
- [11] Yang X., Dove M., Bruin R., Walkingshaw A., Sinclair R., Wilson D.J., Murray-Rust P. (2012) "An e-Science Data Infrastructure for Simulations within Grid Computing Environment: Methods, Approaches, and Practice", to be appeared in "Concurrency and Computation: Practice and Experience", 2012

- [12] Yang X. “*QoS-oriented Service Computing: Bring SOA into Cloud Environment*”, in the Book “Advanced Design Approaches to Emerging Software Systems: Principles, Methodology and Tools”(Edited by Liu X. and Li Y.), Publisher: IGI Global, U.S.A, Publish Date: 2011.
- [13] Zhang S. Wu H., Wang W. Yang B. Liu P. Vasilakos A.V. “Distributed Workload and Response Time Management for Web Applications “, to appear in *IEEE Transactions on Network and Service Management*
- [14] Yang X., Nasser B., Surridge M., Middleton S., “A Business-oriented Cloud Federation Model for Real-time Applications”, *Future Generation Computer Systems*, Elsevier, 2012. Available: <http://dx.doi.org/10.1016/j.future.2012.02.005>. Available: <http://dx.doi.org/10.1016/j.future.2012.02.005>.
- [15] Ludäscher, B., Altintas, I., Berkley, C., Higgins, D., Jaeger, E., Jones, M., Lee, E., Tao, J., Zhao, Y.: “Scientific workflow management and the Kepler system”. *Concurrency and Computation: Practice and Experience*, 18 (10), pp. 1039-1065. 2005.
- [16] Oinn, T., Addis, M., Ferris, J., Marvin, D., Senger, M., Greenwood, M., Carver, T., Glover, K., Pocock, M.R., Wipat, A., Li, P.: Taverna: a tool for the composition and enactment of bioinformatics workflows”. *Bioinformatics*, 20(17), pp. 3045-3054, Oxford University Press, London, UK, 2004.
- [17] Taylor, I., Shields, M., Wang, I., Harrison, A.: The Triana Workflow Environment: Architecture and Applications. In Taylor, I., Deelman, E., Gannon, D., Shields, M. (eds.) *Workflows for e-Science*, pp. 320-339. Springer, New York, Secaucus, NJ, USA, 2007.
- [18] Deelman, E., Mehta, G., Singh, G., Su, M., Vahi, K.: Pegasus: Mapping Large-Scale Workflows to Distributed Resources. In Taylor, I., Deelman, E., Gannon, D., Shields, M. (eds.) *Workflows for e-Science*, pp 376-394. Springer, New York, Secaucus, NJ, USA, 2007.
- [19] Fahringer, T., Jugravu, A., Pillana, S., Prodan, R., Seragiotto, C. Jr., Truong, H.: ASKALON: a tool set for cluster and Grid computing. *Concurrency and Computation: Practice and Experience*, 17(2-4), pp. 143-169, Wiley InterScience, 2005.
- [20] Zhao, Y., Hategan, M., Clifford, B., Foster, I., von Laszewski, G., Nefedova, V., Raicu, I., Stef-Praun, T., Wilde, M.: Swift: Fast, Reliable, Loosely Coupled Parallel Computation. *Proceedings of 2007 IEEE Congress on Services (Services 2007)*, pp. 199-206, 2007.
- [21] Yang, X., Bruin, R., Dove, M (2010) “Developing an End-to-End Scientific Workflow: a Case Study of Using a Reliable, Lightweight, and Comprehensive Workflow Platform in e-Science”, DOI, <http://doi.ieeecomputersociety.org/10.1109/MCSE.2009.211>,
- [22] Ludäscher B, Altintas I, Bowers S, Cummings J, Critchlow T, Deelman E, Roure DD, Freire J, Goble C, Jones M, Klasky S, McPhillips T, Podhorszki N, Silva C, Taylor I, Vouk M (2009) Scientific Process Automation and Workflow Management. In Shoshani A, Rotem D (eds) *Scientific Data Management: Challenges, Existing Technology, and Deployment*, Computational Science Series. 476-508. Chapman & Hall/CRC
- [23] Deelman, E., Gannon, D., Shields, M., Taylor, I.: Workflows and e-Science: An overview of workflow system features and capabilities. *Future Generation Comp. Syst.* 25(5): 528-540, (2009)
- [24] Taylor, I., Deelman, E., Gannon, D., Shields, M. (eds), *Workflows for e-Science*. Springer, New York, Secaucus, NJ, USA, ISBN: 978-1-84628-519-6, (2007)
- [25] Yu, Y., Buyya, R.: A Taxonomy of Workflow Management Systems for Grid Computing. *J. Grid Computing*, 2006 (3):171-200, (2006)
- [26] Wang, J., Korambath, P., Kim, S., Johnson, S., Jin, K., Crawl, D., Altintas, I., Smallen, S., Labate, B., Houk, K.N.: Facilitating E-Science Discovery Using Scientific Workflows on the Grid. In Yang, X., Wang, L., Jie, W. (eds), *Guide to e-Science: Next Generation Scientific Research and Discovery*. ISBN: 978-0-85729-438-8, pp. 353-382. Springer (2011)
- [27] Plale, B., Gannon, D., Reed, D.A., Graves, S.J., Droegemeier, K., Wilhelmson, R., Ramamurthy, M.: Towards Dynamically Adaptive Weather Analysis and Forecasting in LEAD. In *International Conference on Computational Science* (2)(2005) 624-631
- [28] El-Rewini, H., Lewis, T., Ali, H.: *Task Scheduling in Parallel and Distributed Systems*. ISBN: 0-13-099235-6, PTR Prentice Hall, 1994
- [29] MacLennan, Bruce J. *Functional Programming: Practice and Theory*. Addison-Wesley, 1990.
- [30] Wang J., Crawl D., Altintas, I. A Framework for Distributed Data-Parallel Execution in the Kepler Scientific Workflow System. In *Proceedings of 1st International Workshop on Advances in the Kepler Scientific Workflow System and Its Applications at ICCS 2012 Conference*
- [31] Islam M., Huang A., Battisha M., Chiang M., Srinivasan S., Peters C., Neumann A. and Abdelnur A., “Oozie: Towards a Scalable Workflow Management System for Hadoop” In *Proceedings of the 1st international workshop on Scalable Workflow Enactment Engines and Technologies (SWEET'12)*.
- [32] Kosar, T., and Balman, M. “A new paradigm: Data-aware scheduling in grid computing”. *Future Generation Computer Systems*, 25(4), 406-413, 2009.
- [33] D. Yuan, Y. Yang, X. Liu, G. Zhang and J. Chen, “A Data Dependency Based Strategy for Intermediate Data Storage in Scientific Cloud Workflow Systems”. *Concurrency and Computation: Practice and Experience*, Wiley, 24(9): 956-976, 2012.
- [34] Karasavvas, K., Wolstencroft K., Mina E., Cruickshank D., Williams A., De Roure D., Goble C., and Roos, M. “Opening new gateways to workflows for life scientists”, in S. Gesing et al. (Eds.) *HealthGrid Applications and Technologies Meet Science Gateways for Life Sciences*, 2012, IOS Press, pp 131-141.
- [35] Yu, J., Buyya, R., Ramamohanarao, K.: Workflow Scheduling Algorithms for Grid Computing. In Xhafa, F., Abraham, A. (eds), *Metaheuristics for Scheduling in Distributed Computing Environments*. ISBN: 978-3-540-69260-7, pp. 173-214, Springer, 2008.
- [36] Dong, F., Akl, S.: *Scheduling Algorithms for Grid Computing: State of the Art and Open Problems*, Technical Report 2006-504, Queen’s University, 2006.

- [37] Wiczcerek, M., Prodan, R., Fahringer, T.: Scheduling of Scientific Workflows in the ASKALON Grid Environment. SIGMOD Record, Vol. 34, No.3, pp. 56-62, 2005.
- [38] Jianwu Wang, Prakashan Korambath, Ilkay Altintas, Jim Davis, Daniel Crawl. "Workflow as a Service in the Cloud: Architecture and Scheduling Algorithms". In proceeding of International Conference on Computational Science (ICCS 2014), 2014.
- [39] Vazirani, V. V.: Approximation Algorithms. ISBN: 3-540-65367-8, Springer, 2003.
- [40] Morton, T., Pentico, D. W.: Heuristic Scheduling Systems: With Applications to Production Systems and Project Management. John Wiley & Sons. ISBN: 0-471-57819-3, 1993.
- [41] De Roure, D., Goble, C. and Stevens, R. (2009) "The Design and Realisation of the myExperiment Virtual Research Environment for Social Sharing of Workflows". *Future Generation Computer Systems* 25, pp. 561-567.
[doi:10.1016/j.future.2008.06.010](https://doi.org/10.1016/j.future.2008.06.010)
- [42] Gabor Terstysanszky, Tamas Kukla, Tamas Kiss, Peter Kacsuk, Akos Balasko, Zoltan Farkas, "Enabling scientific workflow sharing through coarse-grained interoperability", *Future Generation Computer Systems*, Volume 37, July 2014, Pages 46-59, ISSN 0167-739X, <http://dx.doi.org/10.1016/j.future.2014.02.016>.
- [43] K. Plankensteiner, J. Montagnat, R. Prodan, IWIR: a language enabling portability across grid workflow systems, in: Proceedings of Workshop on Workflows in Support of Large-Scale Science (WORKS'11), Seattle, United States, <http://dx.doi.org/10.1145/2110497.2110509>
- [44] Jianwu Wang, Daniel Crawl, Ilkay Altintas, Weizhong Li. Big Data Applications using Workflows for Data Parallel Computing. Accepted by Computing in Science & Engineering, IEEE
- [45] Bell, G., Hey, T., Szalay, A.: Beyond the Data Deluge. *Science*, 6 March 2009, Vol. 323, No. 5919, pp. 1297-1298, DOI: 10.1126/science.1170411, 2009.
- [46] Dean, J., S. Ghemawat, S.: Mapreduce: Simplified data processing on large clusters, *Communications of the ACM*, vol. 51, no. 1, pp. 107-113, 2008
- [47] Moretti, C., Bui, H., Hollingsworth, K., Rich, B., Flynn, P., Thain, D.: All-pairs: An abstraction for data-intensive computing on campus Grids, *IEEE Transactions on Parallel and Distributed Systems*, vol. 21, pp. 33-46, 2010.
- [48] Gu, Y., Grossman, R.: Sector and sphere: The design and implementation of a high performance data Cloud, *Philosophical Transactions of the Royal Society A*, vol. 367, no. 1897, pp. 2429-2445, June 2009.
- [49] Gropp, W., Lusk, E., Skjellum, A., Using MPI: Portable Parallel Programming with the Message Passing Interface, 2nd ed., ser. Scientific And Engineering Computation Series. Cambridge, MA, USA: MIT Press, 1999.
- [50] Chapman, B., Jost, G., van der Pas, R., Kuck, D.: Using OpenMP : Portable Shared Memory Parallel Programming. Cambridge, MA, USA: The MIT Press, 2007.
- [51] Schatz, M.: Cloudburst: Highly sensitive read mapping with mapreduce, *Bioinformatics*, vol. 25, no. 11, pp. 1363-1369, April 2009.
- [52] Langmead, B., Schatz, M. C., Lin, J., Pop, M., Salzberg, S. L.: Searching for snps with Cloud computing, *Genome Biology*, vol. 10, no. 134, November 2009.
- [53] Kalyanaraman, A., Cannon, W.R., Latt, B., Baxter, D.J.: MapReduce implementation of a hybrid spectral library-database search method for large-scale peptide identification. *Bioinformatics*, Advance online access, 2011.
[doi:10.1093/bioinformatics/btr523](https://doi.org/10.1093/bioinformatics/btr523).
- [54] Wang, J., Crawl, D., Altintas, I.: Kepler + Hadoop : A General Architecture Facilitating Data-Intensive Applications in Scientific Workflow Systems. In Proceedings of the 4th Workshop on Workflows in Support of Large-Scale Science (WORKS09) at Supercomputing 2009 (SC2009) Conference. ACM 2009, ISBN 978-1-60558-717-2. (2009)
- [55] Zhang, C., Sterck, H. D.: CloudWF: A Computational Workflow System for Clouds Based on Hadoop. Proc. of the 1st Int. Conf. on Cloud Computing (CloudCom 2009).
- [56] Fei, X., Lu, S., Lin, C.: A MapReduce-Enabled Scientific Workflow Composition Framework. Proc. of 2009 IEEE Int. Conf. on Web Services (ICWS 2009), pp. 663-670, 2009
- [57] Olston, C., Chiou, G., Chitnis, L., Liu, F., Han, Y., Larsson, M., Neumann, A., Rao, V. B. N., Sankarasubramanian, V., Seth, S., Tian, C., ZiCornell, T., Wang, X.: Nova: Continuous Pig/Hadoop Workflows. ACM SIGMOD 2011 International Conference on Management of Data (Industrial Track), Athens, Greece, June 2011.
- [58] Viana, V., de Oliveira D., Mattoso M.: Towards a Cost Model for Scheduling Scientific Workflows Activities in Cloud Environments, pp.216-219, *IEEE World Congress on Services*, 2011.
- [59] Killapi, H., Sitaridi, E., Tsangaris, M.M., Ioannidis, Y.E.: Schedule optimization for data processing flows on the Cloud. In SIGMOD Conference (2011) 289-300.
- [60] Mateescu, G., Gentzsch, W., Ribbens, C. J.: Hybrid Computing - Where HPC meets Grid and Cloud Computing, *Future Generation Computer Systems*, Volume 27, Issue 5, May 2011, pp. 440-453, ISSN 0167-739X, [10.1016/j.future.2010.11.003](https://doi.org/10.1016/j.future.2010.11.003).
- [61] Manish Parashar, Moustafa AbdelBaky, Ivan Rodero and Aditya Devarakonda, Cloud Paradigms and Practices for Computational and Data-Enabled Science and Engineering, in *Computing in Science & Engineering*, 15, 10-18 (2013), DOI:<http://dx.doi.org/10.1109/MCSE.2013.49>
- [62] Basney J., Gaynor J. An oauth service for issuing certificates to science gateways for teragrid users. *TeraGrid '11*, Salt Lake City, UT., July, 2011
- [63] Pearlman J., Craglia M., Bertrand F., Nativi S., Gaigalas G., Dubois G., Niemeyer S., Fritz S. "EuroGEOSS: an Interdisciplinary Approach to Research and Applications for Forestry, Biodiversity and Drought", Available, <http://www.eurogeoss.eu/Documents/publications%20-%20papers/2011%2034ISRSE%20EuroGEOSS%20Pearlman%20et%20al.pdf>
- [64] Simmhan YL, Plale B, Gannon D (2005). A survey of data provenance in e-Science. SIGMOD Record, 34(3):31-36.
- [65] Illumina Company, HiSeqTM Sequencing Systems, http://www.illumina.com/documents/systems/hiseq/datasheet_hiseq_systems.pdf
- [66] Baker, C.J.O. and Cheung K.-H., eds. *Semantic Web: Revolutionizing Knowledge Discovery in the Life Sciences*. 2006

- [67] Berners-Lee, T.: “*Linked Data – Design Issues*”, W3C, (2009). <http://www.w3.org/DesignIssues/LinkedData.html>
- [68] Shaon A., Woolf A., Crompton S., Boczek R., Rogers W. and Jackson M. (2011). An Open Source Linked Data Framework for Publishing Environmental Data under the UK Location Strategy, Terra Cognita Workshop, The ISWIC 2011 Conference, October 2011. <http://iswc2011.semanticweb.org/fileadmin/iswc/Papers/Workshops/Terra/paper6.pdf>
- [69] Shaon A., Callaghan S., Lawrence B., Matthews B., Osborn T. and Harpham C. (2011). Opening up Climate Research : a linked data approach to publishing data provenance, 7th International Digital Curation Conference (DCC11), Bristol, England, 05-07 Dec 2011. <http://epubs.stfc.ac.uk/work-details?w=60958>
- [70] Bechhofer, S., Ainsworth, J., Bhagat, J., Buchan, I. Couch, P., Cruickshank, D., Delderfield, M., Dunlop, I., Gamble, M., Goble, C., Michaelides, D., Missier, P., Owen, S., Newman, D., De Roure, S. and Sufi S. (2010, December). *Why Linked Data is Not Enough for Scientists*, Proceedings of the 6th IEEE e-Science Conference, Brisbane, Australia.
- [71] Shaon A., Callaghan S., Lawrence B., Matthews B., Osborn T. and Harpham C. (2011). Opening up Climate Research : a linked data approach to publishing data provenance, 7th International Digital Curation Conference (DCC11), Bristol, England, 05-07 Dec 2011. <http://epubs.stfc.ac.uk/work-details?w=60958>
- [72] Sauermann L., Cyganiak R. (2008). *Cool URIs for the Semantic Web*. W3C Interest Group Note 03 December 2008. <http://www.w3.org/TR/cooluris/>
- [73] Zhao J., Goble C. and Stevens R. (2004). *Semantic Web Applications to e-Science in silico Experiments*, in the Proceedings of the 13th international World Wide Web conference on Alternate track papers & posters. <http://www.iw3c2.org/WWW2004/docs/2p284.pdf>
- [74] Haase P., Schmidt M., Schwarte A. (2011). *The Information Workbench as a Self-Service Platform for Linked Data Applications*. In the Proceedings of the Second International Workshop on Consuming Linked Data (COLD2011), Bonn, Germany, October 23, 2011. http://ceur-ws.org/Vol-782/HaaseEtAl_COLD2011.pdf
- [75] Lawrence B., Pepler S., Jones C., Matthews B., and Callaghan S. “Citation and Peer Review of Data: Moving Towards Formal Data Publication”, *International Journal of Digital Curation*, vol. 6, no. 2, 2011. <http://www.ijdc.net/index.php/ijdc/article/view/181/265>
- [76] Callaghan S., Pepler S., Hewer F., Hardaker P., Gadian A. "How to publish data using overlay journals: the OJIMS project", Publication Date: 30-October-2009 , Publication: Ariadne Issue 61, Originating URL: <http://www.ariadne.ac.uk/issue61/callaghan-et-al/> File . Last modified: Thursday, 19-Nov-2009 10:59:06 UTC
- [77] Callaghan S., Hewer F., Pepler S., Hardaker P. and Gadian A., "Overlay Journals and Data Publishing in the Meteorological Sciences ", Publication Date: 30-July-2009 Publication: Ariadne Issue 60 Originating URL: <http://www.ariadne.ac.uk/issue60/callaghan-et-al/> File last modified: Thursday, 30-Jul-2009 15:46:43 UTC
- [78] Earl, T. SOA, Cloud Computing & Semantic Web Technology: Understanding How They Can Work Together. 3rd Annual SOA & Semantic Technology Symposium, 2011, presentation, http://www.afei.org/events/1a03/documents/daytwo_keypm_ertl.pdf
- [79] Foster I. and Kesselman C. (eds), “*The Grid: Blueprint for a New Computing Infrastructure*”, Morgan Kaufmann, July 1998. ISBN 1-55860-475-8.
- [80] Fitzgerald S. “*Grid Information Services for Distributed Resource Sharing*”. The Proceedings of HPDC '01 Proceedings of the 10th IEEE International Symposium on High Performance Distributed Computing
- [81] Laure E. Fisher S.M., Frohner A., Grandi C., Kunszt P. et al “Programming the Grid with gLite”, COMPUTATIONAL METHODS IN SCIENCE AND TECHNOLOGY 12(1), 33-45 (2006)
- [82] Romberg M. “*The UNICORE Grid infrastructure*”. Journal Scientific Programming archive Volume 10 Issue 2, April 2002. IOS Press Amsterdam, The Netherlands.
- [83] Dustdar S., Guo Y., Satzger B., Truong H.L.: *Principles of Elastic Processes*. IEEE Internet Computing 15(5): 66-71 (2011)
- [84] Guo L., Guo Y. and Tian X. “*IC Cloud: A Design Space for Composable Cloud Computing*”. The proceedings of IEEE Cloud Computing 2010, Miami, Florida, USA, 5th-10th, July, 2010.
- [85] Risch M., Altmann J., Guo L., Fleming A., Courcoubetis C. “*The GridEcon Platform: A Business Scenario Testbed for Commercial Cloud Services*”. In Book: Grid Economics and Business Models. LNCS, Volume 5745/2009, Springer Berlin / Heidelberg.
- [86] Toni F., Morge M., et.al . “*The ArguGrid Platform: An Overview*”. In Book: Grid Economics and Business Models. LNCS, Volume 5206/2008, Springer Berlin / Heidelberg.
- [87] Wang, J., Korambath, P., and Altintas, I. *A Physical and Virtual Compute Cluster Resource Load Balancing Approach to Data-Parallel Scientific Workflow Scheduling*. In *Proceedings of IEEE 2011 Fifth International Workshop on Scientific Workflows (SWF 2011)*, at 2011 Congress on Services (Services 2011), pages 212-215
- [88] Schaffer, H. E., Averitt, S. F., Hoit, M. I., Peeler, A., Sills, E. D., and Vouk, M. A. (2009). NCSU's virtual computing lab: a Cloud computing solution. *Computer*, 42(7), 94-97.
- [89] Berriman, G. B., Deelman, E., Juve, G., Rynge, M., and Vöckler, J. S. The application of cloud computing to scientific workflows: a study of cost and performance. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1833), 2013.
- [90] Mell P., Grance T., “The NIST Definition of Cloud Computing”, 2009, www.nist.gov/itl/cloud/upload/cloud-def-v15.pdf.
- [91] Jensen J., Downing R., Waddington S., Hedges M., Zhang J., Knight G. (2011) “Kindura – Federating Data Clouds for Archiving”, *Proc. Int'l Symp. on Grids and Clouds, 2011 (Proceedings of Science)*, to appear.
- [92] Greenwood D., Khajeh-Hosseini A., Smith J., Sommerville I. “The Cloud Adoption Toolkit: Addressing the Challenges of Cloud Adoption in Enterprise”, <http://arxiv.org/pdf/1008.1900>.
- [93] EMC Report: “The Diverse and Exploding Digital Universe” IDC White Paper, March 2008, <http://www.emc.com/collateral/analyst-reports/diverse-exploding-digital-universe.pdf>.

- [94] Hedges M, Hasan A. Blanke T. "Management and preservation of research data with iRODS "(2007) *Proceedings of the ACM first workshop on CyberInfrastructure: information management in e-Science*, (doi: [10.1145/1317353.1317358](https://doi.org/10.1145/1317353.1317358))
- [95] EMC Report: Managing Information Storage: Trends 2011-12, <http://www.emc.com/collateral/emc-perspective/h2159-managing-storage-ep.pdf>.
- [96] Schatz M. C., Langmead B., Salzberg S.L., Cloud Computing and the DNA Data Race, *Nat Biotechnol.* 2010 July; 28(7): 691–693.
- [97] Moore, R.W., Wan, M., Rajasekar, A. (2005) "Storage resource broker; generic software infrastructure for managing globally distributed data", *Proceedings of "Local to Global Data Interoperability - Challenges and Technologies"*, Sardinia, Italy, 20-24 June 2005 (DOI: 10.1109/LGDI.2005.1612467)
- [98] Excel DataScope, Microsoft Research, <http://research.microsoft.com/en-us/projects/exceldatascope>.
- [99] Ruiz-Alvarez A., Humphrey M., "A model and decision procedure for data storage in Cloud computing", *ScienceCloud'11*, June 8, 2011, San Jose, California, USA.
- [100] EPSRC Policy Framework on Research Data (2011), <http://www.legislation.gov.uk/ukpga/2000/36/contents>.
- [101] NERC Data Policy (2011), <http://www.nerc.ac.uk/research/sites/data/policy.asp>.
- [102] Kindura, JISC FSD Programme case study, <http://jiscinfonetcasestudies.pbworks.com/w/page/45197715/Kindura>.
- [103] Rochwerger B., Breitgand D., Levy E., Galis A., Nagin K., Llorente I.M., Montero R., Wolfsthal Y., Elmroth E., Caceres J., Ben-Yehuda M., Emmerich W., Gala F. "The Reservoir model and architecture for open federated Cloud computing", *IBM Journal of Research and Development* Volume: 53 Issue: 4, pages 1-11.
- [104] Chine K., "Open Science in the Cloud: Towards a Universal Platform for Scientific and Statistical Computing", *Handbook of Cloud Computing*, 2010, Part 4, 453-474.
- [105] Loutas N., Peristeras V., Bouras T., Kamateri E., Zeginis D., Tarabanis K., "Towards a Reference Architecture for Semantically Interoperable Clouds", *2010 IEEE Second International Conference on Cloud Computing Technology and Science*, pp. 143-150.
- [106] S. Andreozzi, S. Burke, F. Ehm, L. Field, G. Galang, B. Konya, M. Litmaath, P. Millar, JP Navarro, GLUE Specification v. 2.0 (ANL). March 2009.
- [107] Delbru R., Campinas S., Tummarello G. "Searching Web Data: an Entity Retrieval and High-Performance Indexing Model", In *Journal of Web Semantics*, 2011.
- [108] Vogels, W. (2009). "Eventually consistent". *Communications of the ACM* 52: 40. doi:10.1145/1435417.1435432
- [109] Bizer C., Heath T., Berners-Lee T. "Linked Data - The Story So Far", *Int. J. Semantic Web Inf. Syst.*, 5 (3). 2009, pp. 1-22.
- [110] Paskin N., "Digital Object Identifier (DOI®) System", *Encyclopaedia of Library and Information Sciences*, Third Edition, 1: 1, pp. 1586 — 1592 (ISBN: 978-0-8493-9712-7), 2010, http://www.doi.org/overview/DOI_article_ELIS3.pdf.
- [111] Berners-Lee T., Fielding R., Masinter L., Uniform Resource Identifiers (URI): Generic Syntax. Internet Engineering Task Force (IETF) Request for Comments (RFC) 3986, January 2005. <http://www.ietf.org/rfc/rfc3986.txt>.
- [112] Sollins K., Masinter L., 1994. Functional Requirements for Uniform Resource Names. Internet Engineering Task Force (IETF) Request for Comments (RFC) 1737, December 1994, <http://tools.ietf.org/html/rfc1737>.
- [113] Rochwerger B., Breitgand D., Levy E., Galis A., Nagin K., et al. "The RESERVOIR Model and Architecture for Open Federated Cloud Computing", *IBM Journal of Research & Development*, Volume 53, Number 4, 2009.
- [114] X. Yang, J.D. Blower, L. Bastin, V. Lush, A. Zabala, J. Maso, D. Cornford and P. Diaz and J. Lumsden (2012) "An Integrated View of Data Quality in Earth Observation", *Philosophical Transactions of the Royal Society A*, December 10, 2012, doi: 10.1098/rsta.2012.0072 *Phil. Trans. R. Soc. A* 28 January 2013 vol. 371
- [115] R. Bose, J. Frew, Lineage retrieval for scientific data processing: a survey. *ACM Comput. Surv.* 37, 1 (2005), 1–28.
- [116] K.-K. Muniswamy-Reddy, U. Braun, D. A. Holland, P. Macko, D. Maclean, D. Margo, M. Seltzer, and R. Smogor, "Layering in provenance systems," in *Proc of the USENIX Technical Conf.* USENIX Association, 2009, pp. 129–142.
- [117] K.-K. Muniswamy-Reddy., Macko, P., AND SELTZER, M. I. Making a cloud provenance-aware. In Cheney [6]. J. Cheney, Ed. *First Workshop on the Theory and Practice of Provenance*, February 23, 2009, San Francisco, CA, USA, *Proceedings* (2009), USENIX.
- [118] Nair S.K., Porwal S., Dimitrakos T., Ferrer A.J., Tordsson J., Sharif T., Sheridan C., Rajarajan M., Khan A.U., Towards Secure Cloud Bursting, Brokerage and Aggregation, *2010 Eighth IEEE European Conference on Web Services*, p190-196, <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5693261>.
- [119] Plank G., Burton R.A.B. et al., "Generation of histo-anatomically representative models of the individual heart: tools and application", DOI:10.1098/rsta.2009.0056 *Phil. Trans. R. Soc. A* 13 June 2009 vol. 367 no. 1896 2257-2292.
- [120] He, Q., Zhou, S., Kobler, B., Duffy, D., and McGlynn, T. Case study for running HPC applications in public clouds. In *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing* (pp. 395-401). ACM, 2010
- [121] Bientinesi, P., Iakymchuk, R., and Napper, J. HPC on competitive cloud resources. In *Handbook of Cloud Computing* (pp. 493-516). Springer, 2010.
- [122] Vouk, M. A., Sills, E., and Dreher, P. Integration of High-Performance Computing into Cloud Computing Services. In *Handbook of Cloud Computing*(pp. 255-276). Springer US, 2010
- [123] Chadwick K. et al, "FermiGrid and FermiCloud Update", *International Symposium on Grids and Clouds 2012 (ISGC 2012)*, Taipei, Taiwan, 2012.
- [124] Ahmed, W., Wu, Y. W. A survey on reliability in distributed systems. *Journal of Computer and System Sciences*, Volume 79, Issue 8, December 2013, Pages 1243–1255, doi: <http://dx.doi.org/10.1016/j.jcss.2013.02.006>.
- [125] Dai Y. S., Yang B., Dongarra J., and G. Zhang G., "Cloud service reliability: Modeling and analysis," in *PRDC*, 2009. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.214.143&rep=rep1&type=pdf>.

- [126] Wang C., Wang Q., Ren K., Lou W. Privacy-Preserving Public Auditing for Data Storage Security in Cloud Computing. In INFOCOM, 2010 Proceedings IEEE, doi: <http://dx.doi.org/10.1109/INFCOM.2010.5462173> .
- [127] Rellermeyer, J.S., Bagchi, S. Dependability as a Cloud service - a modular approach. In Dependable Systems and Networks Workshops (DSN-W), 2012 IEEE/IFIP 42nd International Conference. <http://dx.doi.org/10.1109/DSNW.2012.6264688>.
- [128] D.Dahiphale, R. Karve, A. V. Vasilakos, H. Liu, Z. Yu, A. Chhajer, J. Wang, and C. Wang, “An Advanced MapReduce:Cloud MapReduce, Enhancements and Applications”, IEEE Transactions on Network and Services Management, March,vol 11,issue 1, 101-115,2014
- [129] Ikeda R., Park H., and Widom J. “Provenance for Generalized Map and Reduce Workflows”. In Proceedings of CIDR’2011, pages 273–283, 2011.
- [130] Crawl D., Wang J., Altintas, I. “Provenance for MapReduce-based Data-Intensive Workflows”. In Proceedings of the 6th Workshop on Workflows in Support of Large-Scale Science (WORKS11) at Supercomputing 2011 (SC2011) Conference, pp 21-29, 2011.
- [131] Muniswamy-Reddy K, Macko P., and Seltzer M. “Provenance for the cloud”. In Proceedings of the 8th Conference on File and Storage Technologies (FAST’10), the USENIX Association, 2010.
- [132] Foster, I., Zhao, Y., Raicu, I., and Lu, S. Cloud computing and grid computing 360-degree compared. In Grid Computing Environments Workshop, 2008 (GCE’08), pp. 1-10, 2008.