Future Generation Computer Systems 75 (2017) 228–238

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

A characterization of workflow management systems for extreme-scale applications

Rafael Ferreira da Silva [a,*], Rosa Filgueira [b,c], Ilia Pietri [d], Ming Jiang [e], Rizos Sakellariou [f], Ewa Deelman [a]

[a] Information Sciences Institute, University of Southern California, Marina del Rey, CA, USA
[b] British Geological Survey, Lyell Centre, Edinburgh, UK
[c] School of Informatics, University of Edinburgh, Edinburgh, UK
[d] Department of Informatics and Telecommunications, University of Athens, Athens, Greece
[e] Lawrence Livermore National Laboratory, Livermore, CA, USA
[f] School of Computer Science, University of Manchester, Manchester, UK

HIGHLIGHTS

- Design requirements of workflow applications and systems in the extreme-scale.
- Survey and classification of 15 popular workflow engines.
- Research gaps between existing WMS and the desired extreme-scale WMS.

ARTICLE INFO

Article history:
Received 16 May 2016
Received in revised form 11 February 2017
Accepted 15 February 2017
Available online 16 February 2017

Keywords:
Scientific workflows
Workflow management systems
Extreme-scale computing
in situ processing

ABSTRACT

Automation of the execution of computational tasks is at the heart of improving scientific productivity. Over the last years, scientific workflows have been established as an important abstraction that captures data processing and computation of large and complex scientific applications. By allowing scientists to model and express entire data processing steps and their dependencies, workflow management systems relieve scientists from the details of an application and manage its execution on a computational infrastructure. As the resource requirements of today's computational and data science applications that process vast amounts of data keep increasing, there is a compelling case for a new generation of advances in high-performance computing, commonly termed as *extreme-scale computing*, which will bring forth multiple challenges for the design of workflow applications and management systems. This paper presents a novel characterization of workflow management systems using features commonly associated with extreme-scale computing applications. We classify 15 popular workflow management systems in terms of workflow execution models, heterogeneous computing environments, and data access methods. The paper also surveys workflow applications and identifies gaps for future research on the road to extreme-scale workflows and management systems.

1. Introduction

Scientific workflows are an important abstraction for the composition of complex applications in a broad range of domains, such as astronomy, bioinformatics, climate science, and others [1]. Workflows provide automation that increases the productivity of scientists when conducting computation-based studies. Automation enables adaptation to the changing application needs and resource (compute, data, network) behavior. As workflows have been adopted by a number of scientific communities, they are becoming more complex and need more sophisticated workflow management capabilities. A workflow now can analyze terabyte-scale data sets, be composed of a million individual tasks, and can process data streams, files, and data placed in object stores. The computations can be single core workloads, loosely coupled computations (like MapReduce), or tightly coupled (like MPI-based parallel programs) all within a single workflow, and can run in dispersed cyberinfrastructures [1,2].

* Correspondence to: USC Information Sciences Institute, 4676 Admiralty Way Suite 1001, Marina del Rey, CA, 90292, USA.

E-mail addresses: rafsilva@isi.edu (R. Ferreira da Silva), rosa@bgs.ac.uk (R. Filgueira), ipietri@di.uoa.gr (I. Pietri), jiang4@llnl.gov (M. Jiang), rizos@manchester.ac.uk (R. Sakellariou), deelman@isi.edu (E. Deelman).

http://dx.doi.org/10.1016/j.future.2017.02.026
0167-739X/

In recent years, numerous *workflow management systems* (WMSs) have been developed to manage the execution of diverse workflows on heterogeneous computing resources [3–9]. As user communities adopt and evolve WMSs to fit their own needs, many of the features and capabilities that were once common to most WMSs have become too distinct to share across systems. For example, Taverna [8] and Galaxy [9] support advanced graphical user interfaces for workflow composition, making them suitable for bioinformatics researchers with little programming experience. Other systems, such as DAGMan [10] and Pegasus [3] offer scalability, robustness, and planning for heterogeneous high-throughput computation execution. For a new user, choosing the *right* WMS can be problematic simply because there are so many different WMSs and the selection criteria may not be obvious.

To address this problem, several recent surveys [11–18] have been compiled to help users compare and contrast different WMSs based on certain key properties and capabilities of WMSs. These surveys focused mostly on the characterization of the following properties: support for conditional structures (e.g., if and switch statements, while loops, etc.) [12], workflow composition (e.g., graphical user interface, command line, or web portals) [13–16], workflow design (DAG or Non-DAG) [16,17], types of parallelism (e.g., task, data, pipeline, or hybrid parallelism) [14,16,17], computational infrastructure (e.g., cluster, grid, and clouds) [12,14–16], workflow scheduling (e.g., status, job queue, adaptive) [14–18], workflow QoS constraints (e.g., time, cost, reliability, security, etc.) [17], and fault-tolerance and workflow optimizations (e.g., task-level, workflow-level, etc.) [15–17].

Unfortunately, the above characterization properties do not sufficiently address the following question that is on the mind of many computational scientists: "Are WMSs ready to support extreme-scale applications?" We define *extreme-scale applications* as scientific applications that will utilize *extreme-scale computing* to solve vastly more accurate predictive models than before and enable the analysis of massive quantities of data [19,20]. It is expected that the requirements of such applications will exceed the capabilities of current leading-edge high-performance computing (HPC) systems. Examples of extreme-scale applications include: first-principles understanding of the properties of fission and fusion reactions; adaptation to regional climate changes such as sea-level rise, drought and flooding, and severe weather patterns; and innovative designs for cost-effective renewable energy resources such as batteries, catalysts, and biofuels [19].

Extreme-scale computing that includes planned US Department of Energy exascale systems [21] will bring forth multiple challenges for the design of workflow applications and management systems. The next-generation of HPC architectures is shifting away from traditional homogeneous systems to much more heterogeneous ones. Due to the severe energy constraints, data movement will be constrained, both internode and on/off the system, and users will be required to manage deep memory hierarchies and multi-stage storage systems [20,22]. There will be an increased reliance on *in situ* data management, analysis and visualization, occurring in parallel with the simulation [23,24]. These *in situ* processing steps need to be captured to provide context and increase reproducibility.

In addition, as the scientific community prepares for extreme-scale computing, big data analytics is becoming an essential part of the scientific process for insights and discoveries [25]. As big data applications became mainstream in recent years, new systems have been developed to handle the data processing. These include Hadoop [26], a MapReduce-based system for parallel data processing, Apache Spark [27], a system for concurrent processing of heterogeneous data streams, and Apache Storm [28] for real-time streaming data processing. Integrating big data analytics with HPC simulations is a major challenge that requires new workflow management capabilities at the extreme-scale.

In this paper we present a novel characterization of WMSs focused specifically on extreme-scale workflows, using the following properties: (1) workflow execution models, (2) heterogeneous computing environments, and (3) data access methods. Associated with each property is a set of features that can be used to classify a WMS. To evaluate these properties, we select 15 state-of-the-art WMSs based on their broad and active usage in the scientific community, as well as the fact that they have been part of previous surveys. Through a detailed analysis using available publications and other documents, such as project webpages and code manuals, we derive the classification of these WMSs using our features for extreme-scale applications. Our primary contribution in this work is the distillation of all the available information into an easy-to-use lookup table that contains a feature checklist for each WMS. This table represents a snapshot of the state-of-the-art, and we envision it to evolve and grow based on future research in WMSs.

The remainder of this paper is structured as follows. Section 2 presents a background overview of WMSs in general and previous work on characterizing workflows and WMSs. Section 3 describes the two types of extreme-scale workflows that motivate this work. Section 4 presents the three properties for characterizing WMSs for extreme-scale applications, along with their associated features. Section 5 classifies 15 popular WMSs based on these features, and Section 6 describes their current usage in the scientific community. Section 7 identifies gaps and directions for future research on the road to extreme-scale workflows. Finally, Section 8 concludes the paper.

2. Background and related work

2.1. Scientific workflows

The term *workflow* refers to the automation of a process, during which data is processed by different tasks. A WMS aids in the automation of these processes, by managing data and the execution of the application on a computational infrastructure. *Scientific workflows* allow scientists to easily model and express all the data processing tasks and their dependencies, typically as a directed acyclic graph (DAG), whose nodes represent workflow tasks that are linked via dataflow edges, thus prescribing serial or parallel execution of nodes. In general, there are four key properties of scientific workflows, which are handled differently by each WMS:

- *Design:* Most modern WMS provide a graphical user interface to make it easier to create and compose workflows. Alternatively, command line interfaces have the ability to capture more complex structures and scale better to larger problems, but they require programming skills.
- *Execution and Monitoring:* There is a plethora of computing infrastructures, where different optimization methods could be applied to reduce the workflow turnaround time [29]. Robust workflow executions require effective monitoring of a variety of resources including CPU cores, memory, disk, and network traffic [30].
- *Reusability:* WMSs have to make it easier for the workflow designer to reuse previously developed workflows. Many workflows provide mechanisms for tracing provenance and methodologies that foster reproducible science [31].
- *Collaboration:* Due to the collaborative nature of scientific research projects, there is a need for sharing and collaboration mechanisms to foster collaborative efforts among workflow scientists. Some projects, such as myExperiment [32] and Wf4Ever [33], have devoted substantial efforts toward this approach.

In this paper, we are particularly interested in the *Execution and Monitoring* property. We aim to identify, describe, and analyze the different workflow execution models provided by current WMSs to support extreme-scale workflows.