Workflow environments for advanced cyberinfrastructure platforms

Rosa M Badia, Barcelona Supercomputing Center

Progress in science is deeply bound to the effective use of computing infrastructures and to the efficient extraction of knowledge from vast amounts of data. This involves large amounts of data coming from different sources, that follow a cycle composed of pre-processing steps for data curation and preparation for subsequent computing steps, and later analysis and analytics steps applied to results.

In the design of an advanced cyberinfrastructure platform (ACP), one of the key elements is how to describe the applications to be executed in such platform. Most of the times these applications are not standalone, but involve a set of sub-applications or steps composing a workflow. The scientists then rely on effective environments to describe their workflows and engines to manage them in complex infrastructures. From our point of view, current methodologies to describe workflows should be extended in several ways to fulfil the requirements of novel ACPs.

First, most of current approaches focus more on the computing part of the workflow, ignoring the data dimension. Current and, even more, future scenarios will involve large amounts of data coming from different sources (myriad of remote sensors, major scientific instruments, satellites, etc). As described in the BDEC white paper [bdec_paper], while the scientific process involves High-end Data Analysis **(**HDA) steps (abduction and induction), and High-Performance Computing (HPC) steps (deduction), current scientific workflows are performing the three different steps of the scientific process with separated methodologies and tools, with a lack of integration and lack of common view of the whole process. One of the BDEC recommendations is to address the basic problem of the split between the two paradigms: the HPC/HDA software ecosystem split.

We believe that the split between HPC and HDA is due to the fragmentation of the traditional scientific computational workflows into separated components, which use different programming models and different environments for HDA and HPC, with a lack of a global perspective. The huge amount of data and its format heterogeneity, both for the data generated from observations and from deduction, makes very difficult the generation of scientific conclusions. The developers of scientific applications are faced with all this amount of data, large number of data analytics methodologies and HPC tools. Therefore, there is a need for workflow environments and tools for the development of scientific workflows following a holistic approach where both data and computing are integrated in a single flow built on simple, high-level interfaces. Topics of research are novel ways to express the workflows that integrate the different data and compute processes, dynamic runtimes to support the execution of the workflows in complex and heterogeneous computing infrastructures in an efficient way, both in terms of performance and energy.

Besides, the focus of the different scientific and technological user profiles involved in the process may differ. While the emphasis from the computer science point of view has traditionally been on the programming models and applications used to make the predictions and simulations (deduction phase), the scientific application developers give much more emphasis on the data aspect of the problem: metadata that describes the data and traceability of the data that describes how it has been generated or transformed is even more important for them (abduction and induction phases). For this reason, it is common to see scientific workflows environments originated by scientific communities that run very inefficiently (i.e., low CPU utilization) in large HPC systems, while workflow environments that are able to get better performance are not adopted by scientific communities. This can be generalized by differentiating multiple abstraction levels seen by different user profiles that are involved in the use/development of a scientific workflow, from the final user that seeks a higher level of abstraction to define its application, to

the computer engineer that faces herself with a lower level of abstraction with all the details of the infrastructure. In this sense, the design of new workflow environments should follow a multidisciplinary approach, with experts from different Computer Science (CS) fields (machine learning, parallelism, distributed computing) and from application fields involved, in order to define these new methodologies that will support advances in scientific research and knowledge progress. Application providers will contribute to the research together with the CS experts in order to design the workflow environments that better reflect the specific way of understanding their scientific workflows. In this sense, for the different areas of application, different solutions can be designed. What is more, different abstraction levels on the workflow methodologies can be considered to meet the expectations of different scientific and technological user profiles (final user, application developer, research support, computer engineer).

And finally, all these differences and complexities, we need to add the complexity of the current computational infrastructure. We are faced with new processor architectures and of different types (general purpose processors, graphic processors, programmable devices), new persistent storage technologies and new ways of interconnecting all the elements of these complex systems. HPC systems coupled with public and private Cloud infrastructures, and what is more, the systems where future scientific workflows are to be executed will also include edge devices, sensors and scientific instruments that will be able to do computation in the edge and to stream continuous flows of data. Similarly, the scientists expect results to be streamed out for monitoring, steering and visualization of the scientific results to enable interactivity. It is necessary to provide the workflows with powerful runtimes, able to make autonomous decisions in order to execute the scientific workflows in efficient ways in complex data and computing infrastructures, both in terms of performance and energy consumption. The runtime should be able to take decisions in a very dynamic fashion, to enable the exploration of the workflow design space in an intelligent manner, to boost the time to solution. Techniques such as automatic parallelization, machine learning, optimization of data and metadata management, should be present in the runtime. Also, the runtime will be able to deal with the vast and heterogeneous nature of the infrastructures, being able to get the best from them, but keeping the scientific workflow agnostic of them.

[bdec_paper] Asch, M., T. Moore, R. Badia, M. Beck, P. Beckman, T. Bidot, F. Bodin et al. "Big data and extreme-scale computing: Pathways to Convergence-Toward a shaping strategy for a future software and data ecosystem for scientific inquiry." The International Journal of High Performance Computing Applications 32, no. 4 (2018): 435-479.