User Steering Support in Large-scale Workflows

Renan Souza^{1,3} Supervisor: Marta Mattoso¹ Co-supervisor: Patrick Valduriez²

¹Universidade Federal do Rio de Janeiro, Brazil (PESC/COPPE/UFRJ)

²Inria, University of Montpellier, CNRS, LIRMM, France

³Current affiliation: IBM Research

Abstract. Large-scale workflows that execute on High-Performance Computing machines need to be dynamically steered by users. This means that users analyze big data files, assess key performance indicators, fine-tune parameters, and evaluate the tuning impacts while the workflows generate multiple files, which is challenging. If one does not keep track of such interactions (called user steering actions), it may be impossible to understand the consequences of steering actions and to reproduce the results. This thesis proposes a generic approach to enable tracking user steering actions by characterizing, capturing, relating, and analyzing them by leveraging provenance data management concepts. Experiments with real users show that the approach enabled the understanding of the impact of steering actions while incurring negligible overhead.

1. Introduction

How to enable computational scientists and engineers to monitor and understand their workflows when they are steering them on large-scale computers? This is the central question we address in this thesis. Computational Science and Engineering (CSE) workflows are large-scale, require High Performance Computing (HPC) execution, and have the exploratory nature of science [Rüde et al. 2018]. In the long run, which often lasts for hours or days, a typical activity is to have users steering the workflow by dynamically analyzing it (e.g., inspect, monitor) and adapting it (e.g., tune parameters) to improve the quality of results or to reduce the execution time [Souza et al. 2017b].

We call *user steering actions* the individual user interactions performed during steering [Souza et al. 2019b]. Despite the recent breakthroughs in Artificial Intelligence, the complexity of CSE domains makes human knowledge critical and challenging for making decisions during an experiment [Deelman et al. 2017]. These actions need to be tracked, otherwise, users find it harder to understand how and what needs to be steered, they can steer in a misleading way, it can be difficult to explain the results that were consequences of their actions, and it can be impossible to reproduce the results. The importance of supporting the interaction of users with large-scale workflows has been highlighted in several surveys [Deelman et al. 2017, F. da Silva et al. 2017, Atkinson et al. 2017].

In our extensive literature review, we analyze over 450 papers, among which 60 are closely related to this thesis and thus we further investigate in detail. We find that the state-of-the-art in the context of supporting CSE experiments lacks concepts and techniques to allow for tracking steering actions in large-scale workflows. The consequence is that users are left to register the track manually by, for instance, annotating in a sepa-

rate digital spreadsheet or even in a sheet of paper. Thus, the track is often incomplete, isolated from the workflow data, and not stored in a structured way.

Therefore, the general problem this thesis addresses is: "how to allow for tracking user steering actions in running large-scale workflows?". This is hard because to allow for tracking steering actions, steering action data must be managed, which means to capture the steering action data, relate them to the workflow data accordingly, and store in a database ready for online data analysis. There are several challenges associated with managing steering action data, particularly proposing an efficient steering action data capture for running large-scale workflows and building a database with the steering actions related to the workflow data so users can query it online.

To address this problem, we propose *Workflow Steer* (WfSteer), a generic approach to manage user steering action data in large-scale workflows leveraging provenance data. Since provenance data management concepts have been successfully explored to support online workflow data analysis [Silva et al. 2018b], we leverage provenance one step further for the management of user steering action data [Souza et al. 2019b].

WfSteer characterizes and defines user steering action data and operations, and presents system design principles to enable the management of steering action data by capturing, explicitly relating the actions to the rest of the workflow data, and storing these data efficiently. We design and build two instances of this approach: one is a lightweight tool to be plugged into parallel scripts and the other is to be used within a Parallel Workflow Management System (WMS), which are the two typical ways to conduct CSE experiments in HPC. Using real use cases in the Oil and Gas industry, the experiments show that the proposed approach enables users to understand how their actions directly affect the workflow results at runtime and that the system design principles were essential to add negligible overhead to the HPC workflows.

The main contributions of WfSteer are the following. (I) A characterization of user steering actions and steering action data, which are the two main concepts introduced in this thesis, along with their corresponding definitions and how they relate to the workflow data. We also introduced the notion of provenance of steering actions and a corresponding generic data representation following the W3C PROV standard; (II) Distributed systems techniques to manage steering action data when users steer a running workflow, considering the requirements for workflow scripts and WMS; and (III) Design principles to maintain low execution overhead while managing steering action data. Software engineers building workflow systems can adopt these principles to manage steering action data while adding low execution overhead.

2. User Steering Support

WfSteer's goal is to support the *workflow steering lifecycle*, which is given by: (i) analyzing the workflow data online; (ii) choosing what, when, how to adapt; (iii) adapting the workflow data online; and (iv) analyzing the workflow data again to investigate the consequences of the adaptation, and returning to the first step until the end of the execution. In the long run, one steering action is not enough hence requiring several actions.

Every single steering action generates a new kind of data: user steering action data, which also need to be managed as they contain the context that helps to understand

the steering actions and their influence on the workflow. They consist of data informing: when the action happened, why the user decided to act, how the action occurred, which workflow data were analyzed or adapted, what was happening before and after the action, who acted, and the type of the action itself. Since provenance data are so beneficial, we consider that when a user interacts with the workflow execution, the provenance of the generated steering action data must be stored in a database as well. The track of steering actions is the steering action data properly related to the rest of the workflow data.

We distinguish user steering action data management between two aspects: one is to capture steering actions, create the explicit data relationships with the workflow data, and store in a database for online analysis; and the other is to do this while incurring low overhead for the efficient execution performance. WfSteer encompasses both aspects and can be instantiated to scripts and WMSs, regardless of the application domain.

For the former aspect, we conduct an extensive literature review to analyze the state-of-the-art aiming at characterizing steering actions in large-scale workflows and propose a taxonomy in Figure 1(a). Due to the lack of a clear definition of steer action concepts, we propose a formal definition for steering action data and steering operations based on dataflow concepts. Then, we introduce the notion of provenance of steering actions and provide a new data representation, called PROV-DfA (Fig. 2), which extends the W3C PROV standard for steering action provenance [Souza and Mattoso 2018].

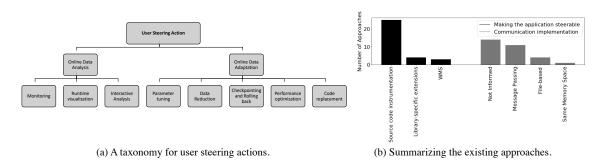


Figure 1. State-of-the-art analysis on user steering actions.

For the latter aspect, we also investigate the way the state-of-the-art approaches support steering actions and summarize this analysis in Figure 1(b). Then, we propose a generic system architecture with a set of design principles to enable tracking the steering actions with low data capture overhead. Some of these principles include a dataflow-oriented data capture [Souza et al. 2017b], an architecture driven by a DBMS that is efficient for both OLAP and OLTP workloads [Souza et al. 2021b], and asynchronous calls to avoid data capture bottlenecks at the running workflows [Souza et al. 2019b]. The sequence diagram in Figure 3 gives an overview of WfSteer services. Based on monitoring the workflow data through a provenance database, users explore intermediate results during a long-lasting execution. Using a provenance data derivation path, users may decide to fine-tune parameters or define subsets of data to focus on.

We also consider the two aspects above to evaluate WfSteer by carrying out two classes of experiments: a qualitative analysis to investigate whether WfSteer allows for tracking steering actions, using real-world workflows as driving use cases; and the other is to quantitatively evaluate the added data capture overhead. Additionally, since Wf-

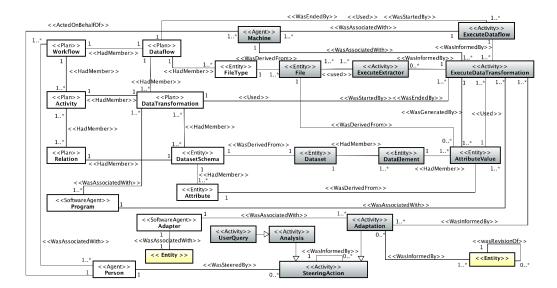


Figure 2. PROV-DfA: a provenance data representation for user steering actions.

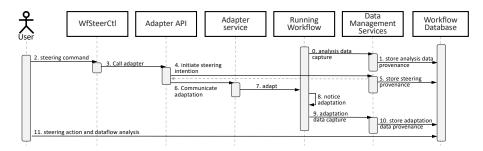


Figure 3. Sequence diagram for managing steering action data.

Steer aims at both scripts and WMSs, within each class of experiments we design and implement two instances of WfSteer: one, called DfAdapter [Souza et al. 2019b]¹, is built on DfAnalyzer [Silva et al. 2018a], to support users who use scripts to conduct their experiments; and the other, implemented within Chiron [Ogasawara et al. 2011, Souza et al. 2015], is to support WMSs' users. DfAnalyzer and Chiron were chosen due to their successful unique features of combining runtime provenance data analysis with negligible overhead in CSE workflows [Camata et al. 2018, Souza et al. 2015].

3. Impact and Conclusions

In this thesis, we aimed at supporting the workflow steering lifecycle to enable computational scientists and engineers to understand their experiments when they are steering workflows. Through the practical use of WfSteer in real cases applied in the Oil & Gas industry, we found that it enabled WMS users to understand how input data were reduced online to yield a reduction of 32.4% of the execution time on an HPC machine with 1,000 CPU cores, hence significantly saving resources [Souza et al. 2017b]; and that in scripts, it allowed users to understand which parameters were tuned that made the workflow finish successfully without memory overflow and with an approximate reduction of 37% (10 days) of the total time using 480 CPU cores of another HPC machine [Souza et al. 2019b].

¹Source code available at https://github.com/hpcdb/DfAdapter.

In these experiments, the user adapted the same parameter multiple times and being able to analyze the past and current context was essential to decide on good values for the next adaptations. We also observed that the formal characterization of steering actions and the adherence to PROV facilitated the modeling of the steering action data relationships.

Finally, we verified that enriching the workflow database with steering action data, jointly with provenance and domain data, enables future interaction analysis. In addition to reliability and reproducibility, we observed from our experiments that having such data enables users to learn from their interactions. For instance, they may find that when they tune specific parameters to a given range of values, the solver convergence improves by a certain amount. Also, AI developers can use these new data to extend training databases of recommendation systems that help users while they are steering simulations. Thus, the computational steering systems should manage steering action data, which has never been done before. We conclude that the user steering action data management concepts and techniques introduced with WfSteer allowed users to track their actions online, enabling them to understand how their steering actions influence a running workflow.

The novel concepts, system architecture design, data representation, design principles, state-of-the-art analyses, and evaluations presented resulted in both journal [Souza et al. 2017b, Souza et al. 2019b, Souza et al. 2021b] conference [Souza et al. 2015, Souza et al. 2016, and Souza et al. 2017a, Souza et al. 2017c, Souza et al. 2018, Souza and Mattoso 2018] papers. The concepts of this thesis influenced collaboration with colleagues [Silva et al. 2016, Silva et al. 2018b, Silva et al. 2018c], as well as Souza's research at IBM [Souza et al. 2021a, Souza et al. 2019a]. We highlight publications in important venues for this research topic, such as Future Generation Computer Systems, International Provenance and Annotation Workshop, Workflows in Support of Large-scale Science, ACM/IEEE Supercomputing, and IEEE e-Science.

References

- Atkinson, M., Gesing, S., Montagnat, J., and Taylor, I. (2017). Scientific workflows: Past, present and future. *Future Generation Computer Systems*, 75:216–227.
- Camata, J. J., Silva, V., Valduriez, P., Mattoso, M., and Coutinho, A. L. G. A. (2018). In situ visualization and data analysis for turbidity currents simulation. *Computers & Geosciences*, 110:23–31.
- Deelman, E., Peterka, T., Altintas, I., Carothers, C. D., Kleese van Dam, K., Moreland, K., Parashar, M., Ramakrishnan, L., Taufer, M., and Vetter, J. (2017). The future of scientific workflows. *International Journal of HPC Applications*, 32(1):159–175.
- F. da Silva, R., Filgueira, R., Pietri, I., Jiang, M., Sakellariou, R., and Deelman, E. (2017). A characterization of workflow management systems for extreme-scale applications. *Future Generation Computer Systems*, 75:228–238.
- Ogasawara, E., Dias, J., Oliveira, D., Porto, F., Valduriez, P., and Mattoso, M. (2011). An algebraic approach for data-centric scientific workflows. *PVLDB*, 4(12):1328–1339.
- Rüde, U., Willcox, K., McInnes, L. C., and Sterck, H. D. (2018). Research and education in computational science and engineering. *SIAM Review*, 60(3):707–754.

- Silva, V., de Oliveira, D., Valduriez, P., and Mattoso, M. (2018a). DfAnalyzer: runtime dataflow analysis of scientific applications using provenance. *PVLDB*, 11(12):2082–2085.
- Silva, V., Neves, L., Souza, R., Coutinho, A., de Oliveira, D., and Mattoso, M. (2018b). Adding domain data to code profiling tools to debug workflow parallel execution. *Future Generation Computer Systems*, 110:422–439.
- Silva, V., Neves, L., Souza, R., Coutinho, A., Oliveira, D. D., and Mattoso, M. (2016). Integrating domain-data steering with code-profiling tools to debug data-intensive workflows. In *Workflows in Support of Large-Scale Science (WORKS) at ACM/IEEE Supercomputing*.
- Silva, V., Souza, R., Camata, J., de Oliveira, D., Valduriez, P., Coutinho, A., and Mattoso, M. (2018c). Capturing provenance for runtime data analysis in computational science and engineering applications. In *International Provenance and Annotation Workshop (IPAW)*, pages 183–187.
- Souza, R., Azevedo, L., Thiago, R., Soares, E., Nery, M., Netto, M., Brazil, E. V., Cerqueira, R., Valduriez, P., and Mattoso, M. (2019a). Efficient runtime capture of multiworkflow data using provenance. In *IEEE e-Science*, pages 1–10.
- Souza, R., G. Azevedo, L., Lourenço, V., Soares, E., Thiago, R., Brandão, R., Civitarese, D., Vital Brazil, E., Moreno, M., Valduriez, P., Mattoso, M., Cerqueira, R., and A. S. Netto, M. (2021a). Workflow provenance in the lifecycle of scientific machine learning. *Concurrency and Computation: Practice and Experience*, pages 1–21.
- Souza, R. and Mattoso, M. (2018). Provenance of dynamic adaptations in user-steered dataflows. In *International Provenance and Annotation Workshop (IPAW)*, pages 16–29.
- Souza, R., Neves, L., Azeredo, L., Luiz, R., Tady, E., Cavalin, P., and Mattoso, M. (2018). Towards a human-in-the-loop library for tracking hyperparameter tuning in deep learning development. In *Latin American Data Science (LaDaS) at VLDB*.
- Souza, R., Silva, V., Camata, J., Coutinho, A., Valduriez, P., and Mattoso, M. (2017a). Tracking of online parameter fine-tuning in scientific workflows. In *Workflows in Support of Large-Scale Science (WORKS) at ACM/IEEE Supercomputing*.
- Souza, R., Silva, V., Camata, J. J., Coutinho, A., Valduriez, P., and Mattoso, M. (2019b). Keeping track of user steering actions in dynamic workflows. *Future Generation Computer Systems*, 99:624–643.
- Souza, R., Silva, V., Coutinho, A., Valduriez, P., and Mattoso, M. (2016). Online input data reduction in scientific workflows. In *Workflows in Support of Large-Scale Science (WORKS) at ACM/IEEE Supercomputing*, pages 1–10.
- Souza, R., Silva, V., Coutinho, A., Valduriez, P., and Mattoso, M. (2017b). Data reduction in scientific workflows using provenance monitoring and user steering. *Future Generation Computer Systems*, 110:481–501.
- Souza, R., Silva, V., Lima, A. A. B., Oliveira, D., Valduriez, P., and Mattoso, M. (2021b). Distributed in-memory data management for workflow executions. *PeerJCS*.
- Souza, R., Silva, V., Miranda, P., Lima, A. A. B., Valduriez, P., and Mattoso, M. (2017c). Spark scalability analysis in a scientific workflow. In *SBBD*, pages 288–293.
- Souza, R., Silva, V., Oliveira, D., Valduriez, P., Lima, A. A. B., and Mattoso, M. (2015). Parallel execution of workflows driven by a distributed database management system. In *ACM/IEEE Supercomputing*, pages 1–3.