



Contents lists available at ScienceDirect

## Future Generation Computer Systems

journal homepage: [www.elsevier.com/locate/fgcs](http://www.elsevier.com/locate/fgcs)

## Raw data queries during data-intensive parallel workflow execution

Vítor Silva<sup>a,\*</sup>, José Leite<sup>a</sup>, José J. Camata<sup>b,c</sup>, Daniel de Oliveira<sup>d</sup>, Alvaro L.G.A. Coutinho<sup>b,c</sup>, Patrick Valduriez<sup>e,f</sup>, Marta Mattoso<sup>a</sup>

<sup>a</sup> Department of Computer Science, COPPE, Federal University of Rio de Janeiro, Brazil

<sup>b</sup> High Performance Computing Center, COPPE, Federal University of Rio de Janeiro, Brazil

<sup>c</sup> Department of Civil Engineering, COPPE, Federal University of Rio de Janeiro, Brazil

<sup>d</sup> Institute of Computing, Fluminense Federal University, Brazil

<sup>e</sup> Inria, France

<sup>f</sup> LIRMM, France

## HIGHLIGHTS

- Workflow algebraic operators to support raw data extraction, indexing and querying.
- Component-based architecture for the analysis of raw data from multiple files.
- Combining the power of runtime provenance data queries with raw data analysis.
- A PROV-Compliant Provenance Data representation to specialize domain data files.
- Real-life parallel execution for Computational Fluid Dynamics workflows.

## ARTICLE INFO

## Article history:

Received 14 May 2016

Received in revised form

15 October 2016

Accepted 7 January 2017

Available online xxxx

## Keywords:

Scientific workflows

Dataflow

Raw data analysis

Index raw data

## ABSTRACT

Computer simulations consume and produce huge amounts of raw data files presented in different formats, e.g., HDF5 in computational fluid dynamics simulations. Users often need to analyze domain-specific data based on related data elements from multiple files during the execution of computer simulations. In a raw data analysis, one should identify regions of interest in the data space and retrieve the content of specific related raw data files. Existing solutions, such as FastBit and RAW, are limited to a single raw data file analysis and can only be used after the execution of computer simulations. Scientific Workflow Management Systems (SWMS) can manage the dataflow of computer simulations and register related raw data files at a provenance database. This paper aims to combine the advantages of a dataflow-aware SWMS and the raw data file analysis techniques to allow for queries on raw data file elements that are related, but reside in separate files. We propose a component-based architecture, named as ARMFUL (Analysis of Raw data from Multiple Files) with raw data extraction and indexing techniques, which allows for a direct access to specific elements or regions of raw data space. ARMFUL innovates by using a SWMS provenance database to add a dataflow access path to raw data files. ARMFUL facilitates the invocation of *ad-hoc* programs and third party tools (e.g., FastBit tool) for raw data analyses. In our experiments, a real parallel computational fluid dynamics is executed, exploring different alternatives of raw data extraction, indexing and analysis.

© 2017 Elsevier B.V. All rights reserved.

## 1. Introduction

Several data-intensive computer simulations take a long time to execute even when using High Performance Computing (HPC)

environments. When these simulations dataflows are managed by a parallel Scientific Workflow Management System (SWMS) [1], they benefit from provenance [2] and data parallelism among different programs that compose the workflow. Systems like Swift/T [3] and Pegasus [4] are highly scalable SWMS and have shown impressive performance results for many different scientific application domains [3,5].

A data analysis challenging problem occurs when users have to navigate and browse thousands of raw data files that result from these data-intensive simulations. Provenance data from SWMS

\* Corresponding author.

E-mail addresses: [silva@cos.ufrj.br](mailto:silva@cos.ufrj.br) (V. Silva), [jvdleite@cos.ufrj.br](mailto:jvdleite@cos.ufrj.br) (J. Leite), [camata@nacad.ufrj.br](mailto:camata@nacad.ufrj.br) (J.J. Camata), [danielcmo@ic.uff.br](mailto:danielcmo@ic.uff.br) (D. de Oliveira), [alvaro@nacad.ufrj.br](mailto:alvaro@nacad.ufrj.br) (A.L.G.A. Coutinho), [patrick.valduriez@inria.fr](mailto:patrick.valduriez@inria.fr) (P. Valduriez), [marta@cos.ufrj.br](mailto:marta@cos.ufrj.br) (M. Mattoso).

<http://dx.doi.org/10.1016/j.future.2017.01.016>

0167-739X/© 2017 Elsevier B.V. All rights reserved.

are an important asset in relating these files, but still very far from supporting raw data analytical queries in these file contents. There are several solutions to improve data analysis through query processing on raw data file contents [6–9]. They typically parse the raw data file; extract relevant contents; index; and present query support, usually with the help of a Database Management System (DBMS). However, these solutions do not support queries that correlate elements from different raw data files.

Combining provenance support from SWMS and raw data query solutions bring a new vision to scientific data analysis. Current solutions are independent and offline [10]. This means that users are only able to query provenance or raw data when the execution finishes. In data-intensive computer simulations, workflow executions may take very long to execute (hours or days) even in HPC environments [5]. Typically, a user tries several different workflow configurations before reaching satisfactory parameters, convergence and error values. This requires a runtime data analysis support, where users may abort or fine tune, and debug their workflow long before it finishes.

In a previous work [11] we have shown the advantages of adapting loop conditions based on partial data analysis, all during the iterative workflow execution. We used Chiron SWMS [12] and its data-centric algebra [13] to query provenance data related to domain data using a specialization of W3C PROV [14]. The work in [11] evolved into defining workflow algebraic operations to control and adapt loop conditions [15]. In a study on uncertainty quantification we anticipated a convergent state and dynamically changed loop control [15]. Such actions contributed to reduce execution time in several hours. These simulations or other iterative workflows typically take several hours or days to execute. These results led us to improve raw data analysis support into selecting raw data elements from files and relating them to workflow parameters at runtime [16]. However, our previous solutions suffer from limited raw data access with no direct path to specific regions or elements of raw data files. Users still had to write specific programs to access and analyze the raw data files. Even though direct access to files can be obtained by querying the provenance database, the raw data file content analysis remains isolated from the provenance database.

There are several challenges in providing raw data file analysis, while the workflow is being executed by the SWMS. Performance is critical; the overhead in runtime data analysis support must not harm the parallel workflow execution time. Another challenge is managing the large size of these files with their specific raw data format; they cannot be converted to be inserted (*i.e.*, replicated) into a database for queries. There is also the issue of mapping and accessing specific regions of interest inside the raw data file.

To address these challenges and benefit from SWMS provenance data with raw data analysis support, the following services must be provided during the workflow execution:

- access to raw data files while they are being generated;
- parse raw data to find relevant data;
- extract relevant subsets of raw data;
- index over data regions of interest;
- prepare raw data for queries;
- runtime query relating raw data from different files, provenance data, and performance execution data.

In this paper, we present a raw data analysis support to address these challenges. We implemented all these services in an architecture named ARMFUL (Analysis of Raw data from Multiple Files) that can be plugged in SWMS. In its current version, it is implemented in an extended version of Chiron and evaluated in an HPC environment with a real numerical simulation workflow. Experiments with a finite element solver for fluid dynamics [17] workflow show the performance improvements in

runtime queries obtained by the ARMFUL indexing techniques. The results present relevant costs/benefits considering the overheads of managing and indexing raw data, with further gains from powerful runtime data analyses obtained by queries accessing data directly.

The rest of the paper is organized as follows. Section 2 describes a motivating scenario from numerical simulations. Section 3 discusses related work on raw data analysis from the execution of computer simulations. Section 4 defines dataflow concepts with workflow algebra operations to analyze raw data files. Section 5 presents ARMFUL, an architecture to support raw data analysis, which is based on raw data extraction and index generation. Section 5 also presents how we use Chiron SWMS to implement ARMFUL. Section 6 uses the motivating example from Section 2 to present the experimental results. ARMFUL raw data support is evaluated with real computational fluid dynamics workflows, while extracting, indexing and querying raw data from XDMF and HDF5 files. Finally, Section 7 concludes.

## 2. Motivating scenario

To clarify the exploratory analysis of raw data files, let us consider an example from Computational Fluid Dynamics (CFD) that is presented in Fig. 1. Specifically, this simulation analyzes the three-dimensional flow of an incompressible fluid in a cavity. This problem is a popular benchmark in CFD, used to evaluate new codes or new solution methods [18]. This benchmark is consistently used throughout this paper. In Fig. 1, the black boxes show the simulation programs invoked for the configuration of the CFD analysis (*edgcfdfPre* program) and solver execution (*edgcfdfSolver* program). The *edgcfdfPre* program generates configuration files in *part.in* and *part.mat* formats, which contain the needed raw data elements, such as the maximum time step (attribute *DTMAX*), the final simulation time (attribute *TMAX*), material properties such as fluid viscosity (attribute *VISC*) and fluid density (attribute *DENS*). The *edgcfdfSolver* program consumes the produced files by *edgcfdfPre* program, and generates solution attributes (such as velocity and pressure) over the execution of the CFD solver (*i.e.*, files in XDMF and HDF5 formats) and some metrics for analyzing the CFD model convergence (number of linear iterations and residual norms represented as attributes *ITER* and *RESIDUALS*, in STP file format). The attribute *RESIDUALS* represents the ratio of the Euclidean norm of the momentum and continuity residuals between two linear iterations.

Users in CFD domain commonly investigate the convergence of their solver based on the number of linear iterations and residual norms for a specific input mesh, boundary, initial conditions and material properties. To access solution quality and/or to define if it is necessary more time steps, users may increase the value of *TMAX*. This means that the CFD solver needs more time to reach a steady state, and other controlling parameters may also be adjusted. To analyze the number of linear iterations and residual norms, users must gather raw data elements (*e.g.*, values for *ITER* and *RESIDUALS*) from STP files. To narrow their analysis for a specific input mesh and fluid properties and correlate it to the solver properties, users have to investigate the dataflow path for a specific input mesh (*e.g.*, *cav.1.msh*, *cav.1.part.in*, *cav.1.xmf*, and *cav.1.stp* files). Furthermore, checking if the solution has reached the steady state (solution is no longer changing) is based on the reduction on the number of linear iterations and residual norms over time, as presented by the behavior in charts of scenario (a) in Fig. 1. A solver execution that is not able to converge is represented by the behavior of scenario (b) in Fig. 1. The analysis of scenario (b) can assist users to abort CFD executions that will not converge or to change some attributes like *TMAX*. The modifications on *TMAX* need a deep analysis in

Download English Version:

<https://daneshyari.com/en/article/4950427>

Download Persian Version:

<https://daneshyari.com/article/4950427>

[Daneshyari.com](https://daneshyari.com)