Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 108C (2017) 2190–2199

Procedia Computer Science

International Conference on Computational Science, ICCS 2017, 12-14 June 2017, Zurich, Switzerland

Design Evaluation of a Performance Analysis Trace Repository

Richard Grunzke[1]*, Maximilian Neumann[1], Thomas Ilsche[1], Volker Hartmann[2], Thomas Jejkal[2], Rainer Stotzka[2], Andreas Knüpfer[1], and Wolfgang E. Nagel[1]

[1] Center for Information Services and High Performance Computing
Technische Universität Dresden, Dresden, Germany
richard.grunzke@tu-dresden.de
[2] Institute for Data Processing and Electronics
Karlsruhe Institute of Technology, Karlsruhe, Germany

Abstract

Parallel and high performance computing experts are obsessed with performance and scalability. Performance analysis and tuning are important and complex but there are a number of software tools to support this. One methodology is the detailed recording of parallel runtime behavior in event traces and their subsequent analysis. This regularly produces very large data sets with their own challenges for handling and data management. This paper evaluates the utilization of the MASi research data management service as a trace repository to store, manage, and find traces in an efficient and usable way. First, we give an introduction to trace technologies in general, metadata in OTF2 traces specifically, and the MASi research data management service. Then, the trace repository is described with its potential for both performance analysts and parallel tool developers, followed with how we implemented it using existing metadata and how it can utilized. Finally, we give an outlook on how we plan to put the repository into productive use for the benefit of researchers using traces.

© 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the scientific committee of the International Conference on Computational Science

Keywords: Performance Analysis, Data Repository, Metadata Extraction

1 Introduction

Many fields of science and research today are driven by data. The amount of data, the numbers of files, and the complexity of research data are growing steadily. Furthermore, data is not assigned to individuals anymore but teams are required to work with it in a joint manner. Due to this, organized data and metadata management is indispensable. It is the basis for collaborative work with data sets and it provides valuable benefits like 1) automatic annotation of advanced metadata; 2) advanced search capabilities for finding pieces of data; 3) interoperable (high performance) data access and sharing; 4) integration of standard data (pre-)processing; and 5) support for data preservation/archiving.

1877-0509 © 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the scientific committee of the International Conference on Computational Science
10.1016/j.procs.2017.05.190

In this paper we consider the scenario of trace-based performance analysis for parallel high performance computing. This method is based on a very detailed recording of the parallel execution of a target application which is then subject to post-mortem analysis. The findings of this analysis then guide the performance tuning process. Depending on various factors the involved trace files easily reach into the 3-digit gigabyte range for one execution of a parallel application. Usually, several traces need to be recorded in the course of one tuning project. Furthermore, the history from many such tuning projects is a valuable resource to learn from for future projects. Thus, this paper presents a very data intensive scenario with its own challenges for data and metadata management. We implemented the first version of a repository that is able to manage event trace files in a structured way. Researchers can search for traces according to various characteristics. This is a substantial shift, as before the traces were stored in a shared directory structure and users had to laboriously search for traces that meet their requirements in a manual way. Now, users can graphically search for the specific characteristics of traces based on structured metadata.

First, Section 2 describes the relevant trace technologies and the MASi service. Then, Section 3 details the use case, the implementation, and its usage. The move towards the sustainable operation is outlined in Section 4 while in Section 5 related work is presented.

2 Background

2.1 Trace Technologies

Performance analysis techniques can be characterized by a number of aspects. In this work, we look in particular at post-mortem performance analysis, where measurement data is collected during the experiment and written to persistent storage for later analysis. Further, we focus on tracing tools that log each individual measurement event during data recording, as opposed to profiling tools that summarize the collected data. Tracing results in a large amount of data, therefore it is particularly important to use metadata to improve the accessibility of performance measurements. A more detailed description of the classification of performance measurement tools can be found in [7].

There are several tools for performance analysis of HPC applications, most of them using their own measurement format. The Tuning and Analysis Utilities (TAU) [14] focus mostly on profiling, but also provide some support for tracing. HPCToolkit [15] measures the application using sampling and can record either profiles or traces. The Open Trace Format (OTF) [10] presents an effort towards a trace format that can be used independently of a specific tool. It was mainly developed as an output format for the VampirTrace [13] measurement tool. Similarly, its successor OTF2 [3] is associated with the Score-P measurement infrastructure [12]. Score-P and the resulting OTF2 measurements are used by the analysis and visualization tools Scalasca [4] and Vampir [13]. These tool chains can tackle parallel traces of up to several terabytes.

The data of a performance trace consists of a series of time-stamped events or samples that can have many different attributes. This variety of events in the trace data, already presents a complex opportunity for filtering and searching multiple traces. In addition, a wide range of metadata can be associated with a trace.

2.2 Metadata in OTF2 Traces

OTF2 traces created by Score-P contain various metadata:

دریافت فوری ← متن کامل مقاله

ISIArticles
مرجع مقالات تخصصی ایران

✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
✓ امکان دانلود نسخه ترجمه شده مقالات
✓ پذیرش سفارش ترجمه تخصصی
✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
✓ دانلود فوری مقاله پس از پرداخت آنلاین
✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات