

Large-Scale Trace Visualization Analysis with Triva and Pajé – the G5K case study –

Lucas M. Schnorr, Arnaud Legrand

Grid'5000 Spring School
Reims, France

20 April 2011

Outline

- 1 Introduction
- 2 Basic Concepts
- 3 Visualization Technique
- 4 Case Studies

Introduction

- Performance Analysis may lead to
 - Improve the code of the parallel application
 - Deploy processes according to platform
- Register the behavior of processes
 - Profiling: gives an overview of where time is spent
 - Tracing: detailed behavior with timestamped events
- Visual Analysis and Understanding of Traces

Objectives of the Practical Session

- Time & Spatial Integration for scalable analysis
- Get acquainted with the Triva Visualization Tool
- Locate anomalies and pin-point network bottleneck

Pajé File Format

- Generic text-based file format, without semantics
- Pajé Types (all have be associated to a container)
 - **container**, represents the monitored entity
 - **event**, represents events at a certain instant
 - **state**, given state for a given period of time
 - **link**, relation between two containers with start/end
 - **variable**, evolution in time of a certain value
- Each file has three sections
 - 1 Header (events declaration, lines starting with %)
 - 2 Types Definition (organized as a type hierarchy)
 - 3 Timestamped events (starting at time 0.0 seconds)

Visualization Scalability

- Parallel or Distributed Applications
 - Time: Many traceable events for each process
 - Space: several thousands processes
- Analysis has to take into account both dimensions
- Not feasible to look into every single entity/event

What are our options to scale?

- Time Integration
- Spatial Integration

Time Integration

- Reduce details in the time dimension
- Considering a time-slice defined by the analyst
 - **event** type is counted
 - the value of **state**, **variable** types are time integrated
- The time-slice definition can be dynamic
 - The analyst can change it interactively
 - Forward and backward according to the analysis
 - The visual representation changes accordingly

Spatial Integration

- Reduce details in the space dimension
- Based on a distance metric to group entities
- In Pajé, the distance is the hierarchy of the trace
 - **containers** of one type are merged into the parent
 - can be recursively applied
- Example of a Pajé type hierarchy with
 - Site – Cluster – Machine – Process
 - All processes of a machine can be grouped together
 - All machines of a cluster
 - All clusters of a site
 - All sites

Treemap View

- A complementary technique to analyze traces
 - Highlight differences in processes behavior
 - Adapted to explore space-integrated data
 - Scalable analysis to several thousands entities
- Hierarchy Nodes' values associated to screen space
- How it works
 - Time/Space-integrated metrics from traces are used to define the hierarchy nodes' value
 - Relate screen space to a given trace metric
- Examples of its use:
 - Processes with more time spent in a given state
 - Verify load-balancing between clusters/sites

Graph View

- A complementary technique to analyze traces
 - Topological-based visualization using a graph
 - Application behavior defines the layout of nodes
- Detection of network bottlenecks, in time/space
- How does it work
 - **containers** are nodes, **links** are edges
 - **events, variables, states** are used to define layout
- Layout: shape (square or rhombus), size, utilization
- Position: Force-based algorithm
 - with interactive spring/charge/damping parameters

Practical Session

- Treemap-based Analysis
 - KAAPI Work Stealing traces (Fibonacci Application)
 - 1 Fibonacci application, 188 processes
 - 5 sites, 188 machines allocation
 - Rennes, Nancy, Toulouse, Bordeaux, Porto Alegre

- Graph-based Analysis
 - Traces from the NAS Data Traffic (DT) Benchmark
 - 1 DT Class A, Black Hole, 21 processes
 - 2 DT Class C, White Hole, 85 processes
 - 3 DT Class C, Shuffle, 448 processes
 - 92 machines of Griffon Cluster of Nancy's site, 3 racks

Installing Triva...

- 1 Deploy **triva-squeeze-x64**, available at Nancy
 - Access with `ssh -X` (g5k:grid5000)
 - (not so good because of network latency)
- 2 Instal Triva manually (Linux and Mac)
 - `svn checkout svn://scm.gforge.inria.fr/svn/triva`
 - `INSTALL` file
- 3 Install Triva automatically (room machines)
 - `wget http://triva.gforge.inria.fr/g5k-triva.sh`
 - `source gk-triva.sh`
 - `Triva -help`

References

- Triva – <http://triva.gforge.inria.fr>
- Pajé – <http://paje.sourceforge.net>