



<https://hps.vi4io.org>

Zoya Masih

Trace Analysis in Practice

Outline

- 1 Introduction
- 2 Score-P
- 3 Vampir

Table of Contents

1 Introduction

2 Score-P

3 Vampir

Motivation

- In the previous sessions we introduced:
 - ▶ profiling as aggregated statistics
 - ▶ tracing as a detailed event timeline
- Score-P combines both ideas in one HPC toolchain
- Vampir is one of the main interfaces to explore Score-P traces

Objectives

Goal of this session

- Understand the workflow
- Run it on the cluster
- Interpret the output

Score-P and Vampir: Big Picture

- **Score-P** collects performance data from HPC applications
- **Vampir** is commonly used for trace analysis

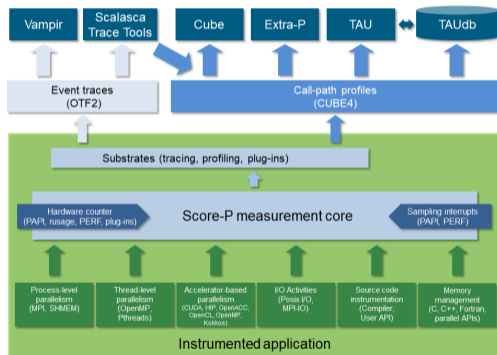


Figure: <https://helmholtz.software>

Table of Contents

1 Introduction

2 Score-P

3 Vampir



Measurement Procedure

- Instrumentation inserts measurement probes into the code
- Probes are triggered during execution
- Execution records events and performance data
- Data includes time, visits, communication , hardware counters
- Output is stored as profile or trace for analysis

Output Files

- **scorep.cfg:** Listing of used environment variables
- **profile.cubex:** CUBE4 result file of the summary measurement
- **traces.otf2:** the anchor file for the OTF2 (Open Trace Format 2) trace, used to capture and storing performance data. the main entry point for accessing trace data.
- **traces.def:** contains global definitions for the trace data. Includes info about the structure and organization of the trace, like the definitions of events, locations, regions, etc.
- **traces/:** containing per-location trace data. It contains detailed trace info specific to each execution location (e.g., process or thread).

Analyzing Results without GUI Tools


- Score-P produces useful output even without GUI tools
- `scorep-score profile.cubex` provides:
 - ▶ time, visits, and region statistics
- Additional processing (advanced):
 - ▶ remapping can compute extra metrics and organize data hierarchically

```
Estimated aggregate size of event trace:          334 bytes
Estimated requirements for largest trace buffer (max_buf): 334 bytes
Estimated memory requirements (SCOREP_TOTAL_MEMORY): 4097kB
(hint: When tracing set SCOREP_TOTAL_MEMORY=4097kB to avoid intermedia
or reduce requirements using USR regions filters.)
```

flt	type	max_buf[B]	visits	time[s]	time[%]	time/visit[us]	region
	ALL	333	9	4.64	100.0	515655.50	ALL
	MPI	240	6	3.33	71.7	554434.59	MPI
	SCOREP	41	1	0.00	0.0	33.63	SCOREP
	USR	26	1	1.31	28.2	1308490.36	USR
	COM	26	1	0.01	0.1	5767.92	COM

Table of Contents

- 1 Introduction
- 2 Score-P
- 3 Vampir**



How to use Vampir

- Vampir on a Login Node
- Vampir Using VampirServer
- Vampir on Local Machine (limited)





Figure: docs.olcf.ornl.gov

Process Timeline

- Information about different levels of function calls
- Levels show the execution flow and hierarchical relationships




Process Timeline

- Shows when each rank or thread is active, waiting, or inside a specific function



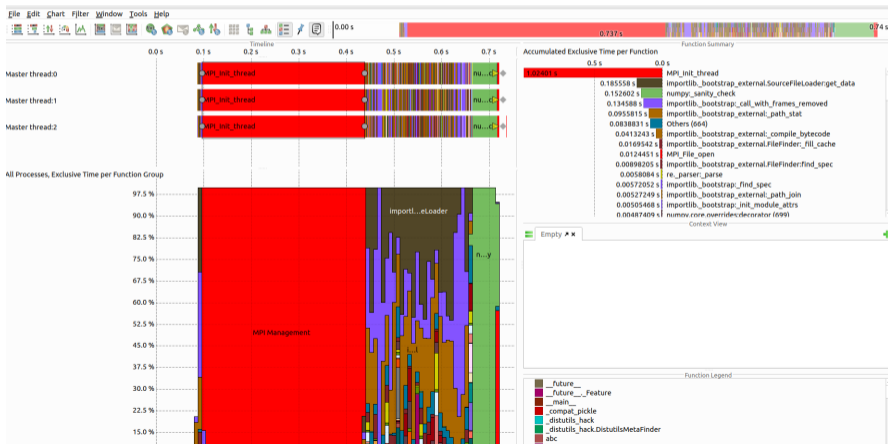
Summary Timeline

- Provides detailed information about functions, communications, and synchronization




Summary Timeline

- Helps identify dominant phases and whether behavior is balanced across processes



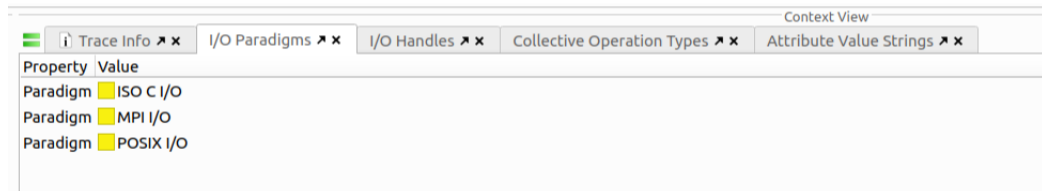
I/O summary

- Vampir highlights I/O operations if I/O performance data has been recorded.
 - ▶ `-io=runtime:posix`
- MPI I/O routines are automatically instrumented
 - ▶ if the scorep instrumenter detects the MPI programming paradigm



I/O summary

- Confirms which I/O interfaces were recorded, such as POSIX or MPI I/O



The screenshot shows the Vampir I/O summary window. At the top right, it says "Context View". Below that are several tabs: "Trace Info", "I/O Paradigms", "I/O Handles", "Collective Operation Types", and "Attribute Value Strings". The "I/O Paradigms" tab is selected. Below the tabs is a table with two columns: "Property" and "Value".

Property	Value
Paradigm	ISO C I/O
Paradigm	MPI I/O
Paradigm	POSIX I/O

Hands-on Plan

- **Exercise 1:** Build and trace a small C program
- **Exercise 2:** Analyze the output traces in Vampir
- **Exercise 3:** Compare the results with an optimized version
- **Exercise 4:** Build and trace a small Python program
- **Exercise 5:** Analyze the output traces in Vampir