



Alexander Movsesyan, Rakrish Dhakal, Aditya Ranjan, Jordan Marry, Onur Cankur, Abhinav Bhatele **Department of Computer Science, University of Maryland**

Abstract

Performance analysis is an imperative part of performance tuning during the development of parallel programs. Parallel execution traces enable in-depth analysis of the program's performance. Current trace analysis tools/workflows have some gaps:

- Most trace analysis tools support different formats and analyses
- GUI-based tools limit data exploration to their graphical views

We have developed Pipit, a Python-based tool, to fill in the gaps in trace-analysis:

- Supports traces in different file formats (OTF2, HPCToolkit, Projections, etc.)
- Provides a uniform data structure in the form of a pandas DataFrame
- Provides a programmatic API to analyze traces
- Provides interactive visual functions to display the traces



Background and Pipit's Structure

Traces: Time series data representing all the events that occur during the program's execution

- When functions are entered and exited
- When messages are exchanged between processes
- Different performance metrics (such as hardware performance counters)

How does Pipit store trace data?

- A pandas DataFrame: two-dimensional labeled table-like data structure
- Every trace event is stored as a row in the DataFrame
- DataFrame is sorted by event timestamps

The Calling Context Tree

- Represents caller-callee relationships between functions
- Stored as a graph in Pipit, and each event in the DataFrame corresponds to some node in the calling context tree

imesta ame, P , Ente , Ente , Ente , Leav	amp Prov er, er, er, er,	<pre>(s), Event Ty cess main(), 0 foo(), 0 MPI_Send, 0 MPI_Send, 0 hoz() 0</pre>	pe,		
8, Le		Timestamp (ns)	Event Type	Name	Process
5, Le 00, L	0	0	Enter	main()	0
	1	100000000	Enter	foo()	0
	2	300000000	Enter	MPI_Send	0
	3	500000000	Leave	MPI_Send	0
	4	800000000	Enter	baz()	0
	5	1800000000	Leave	baz()	0
	6	25000000000	Leave	foo()	0
	7	10000000000	Leave	main()	0



Pipit: Simplifying Parallel Trace Analysis



1. Finding the Most Idle Processes

We analyze traces of a Loimos (a Charm++-based epidemiology simulator) execution on 64 processes. We want to find which processes are idling the most while the others are overloaded. Pipit's idle_time function can help us with this task. We then plot a timeline filtered to the most and least idling processes.



3. Pattern Detection

We can use pipit's *detect_pattern* function to find recurring sets of events in the trace. Below, we analyze a Tortuga execution on 16 cores. The function uses matrix profile to detect patterns in the trace. We use this function to automatically identify loops in the program.



tortuga_16 = pipit.Trace.from_otf2('./tortuga_16') matches = tortuga_16.detect_pattern(window_size , iterations , metric='time.exc') tortuga_16.plot_timeline()

References





Performance of Pipit

All the experiments in this section were performed on a single node of an HPC cluster with a dual 64-core AMD EPYC 7763 processor.

Time spent by the Pipit OTF2 reader in reading traces of two different applications, AMG (128 processes) and Laghos (256 processes).



Time spent in the OTF2 reader and the comm_matrix function with AMG and Laghos traces of different sizes.



Getting Started

Install pipit with pip

pip install pipit

Scan the QR code for pipit on GitHub



H Wes McKinney. 2017. Python for Data Analysis: Data Wrangling with Pandas, NumPy, and IPython. O'Reilly M