High Performance Computing Systems (CMSC714)

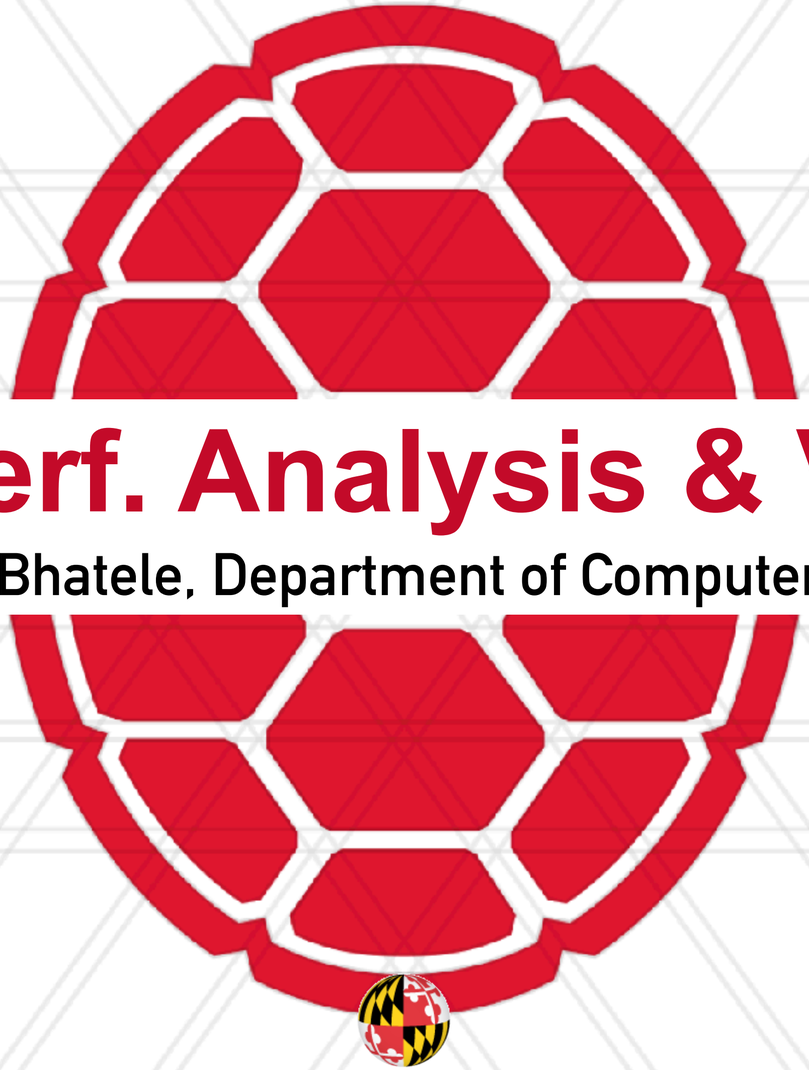

Lecture 8: Perf. Analysis & Visualization

Abhinav Bhatele, Department of Computer Science

UNIVERSITY OF MARYLAND

Summary of last lecture

- Task-based programming models and Charm++
- Key principles:
  - Over-decomposition, virtualization
  - Message-driven execution
- Automatic load balancing, checkpointing, fault tolerance

Tracing tools

- Record all the events in the program with timestamps
- Events: function calls, MPI events, etc.

Vampir visualization: https://hpc.llnl.gov/software/development-environment-software/vampir-vampir-server

Tracing tools

- Record all the events in the program with timestamps
- Events: function calls, MPI events, etc.

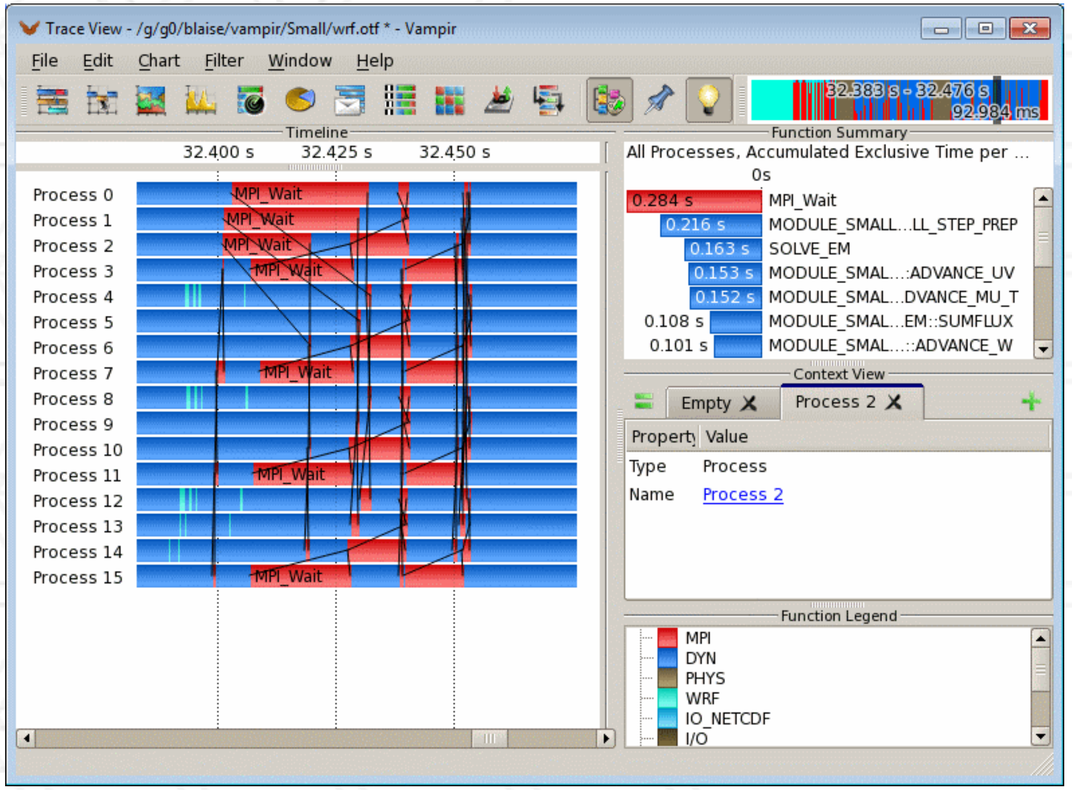

Vampir visualization: https://hpc.llnl.gov/software/development-environment-software/vampir-vampir-server

Tracing tools

- Record all the events in the program with timestamps
- Events: function calls, MPI events, etc.

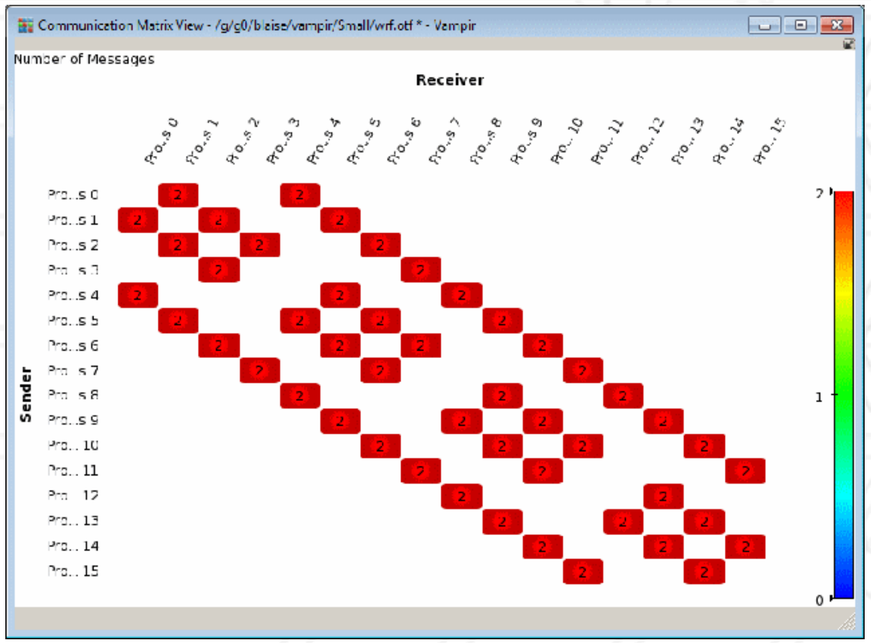

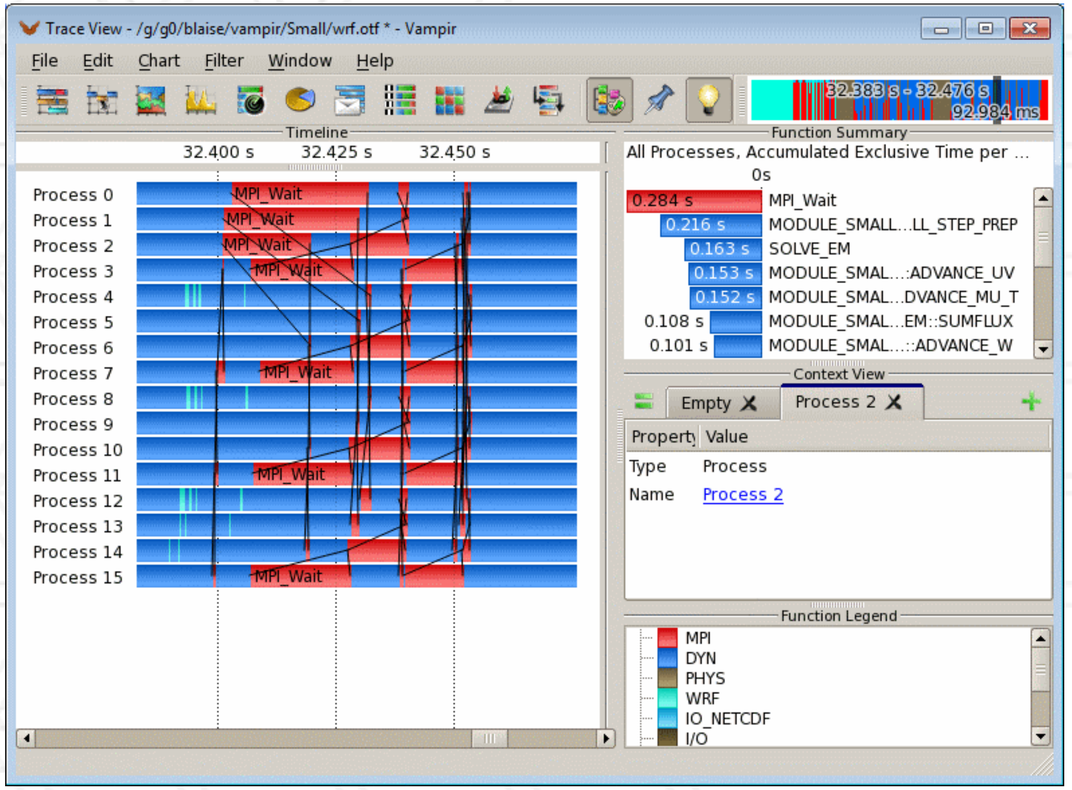

Vampir visualization: https://hpc.llnl.gov/software/development-environment-software/vampir-vampir-server

MPI trace visualization

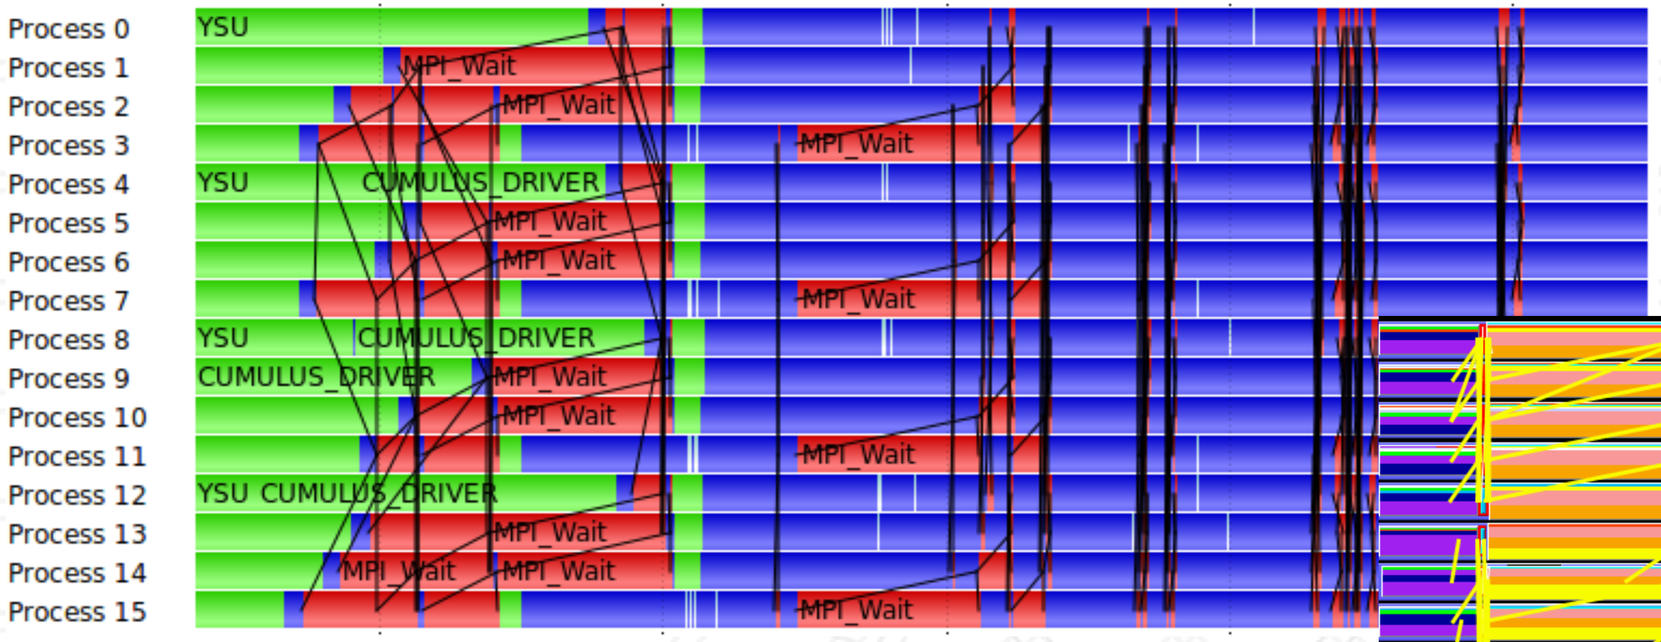

Vampir

Jumpshot

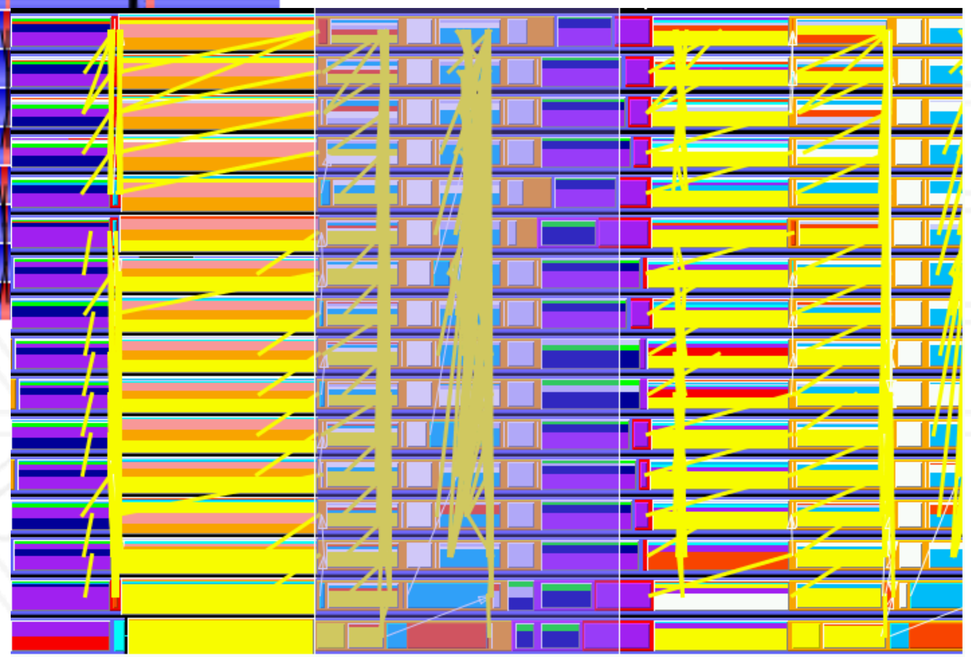

Projections Performance Analysis Tool

- For Charm++/Adaptive MPI programs
- Instrumentation library
  - Records data at the granularity of chares (Charm++ objects)
- Java-based GUI

Time Profile

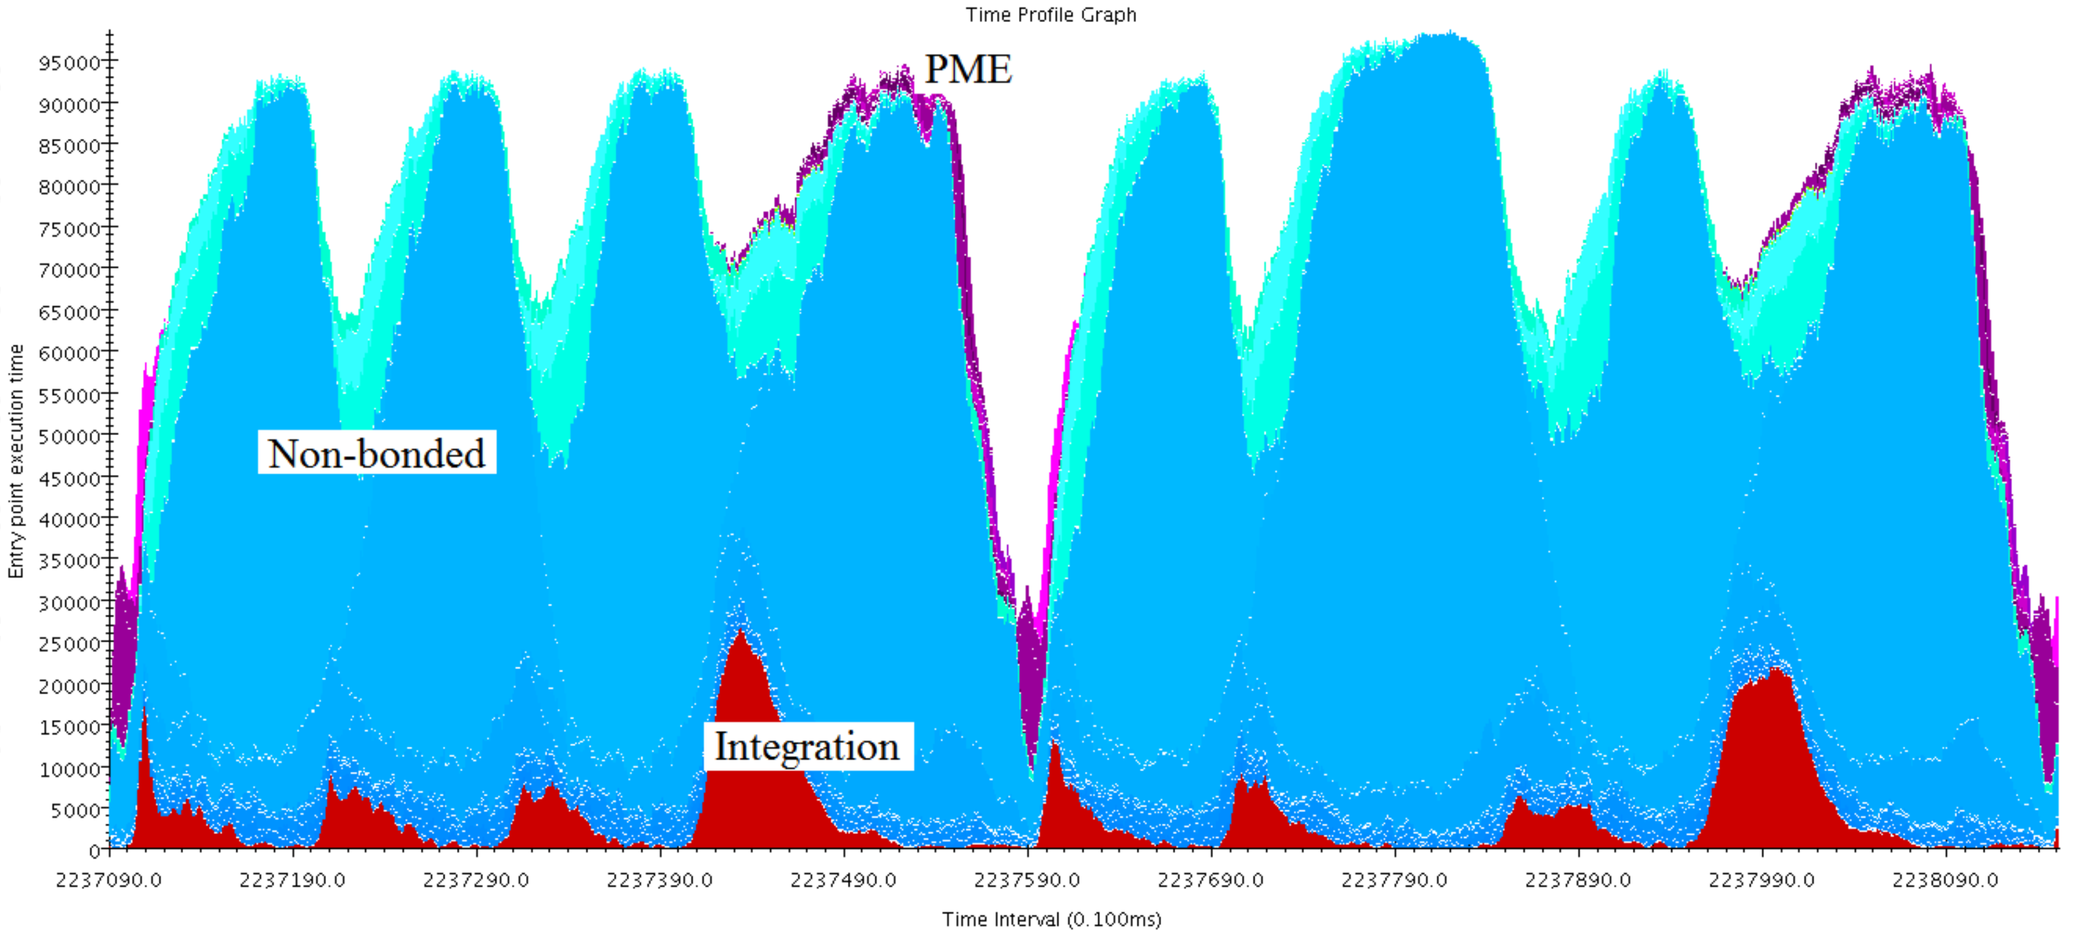

https://charm.readthedocs.io/en/latest/projections/manual.html

Usage Profile & Histogram View

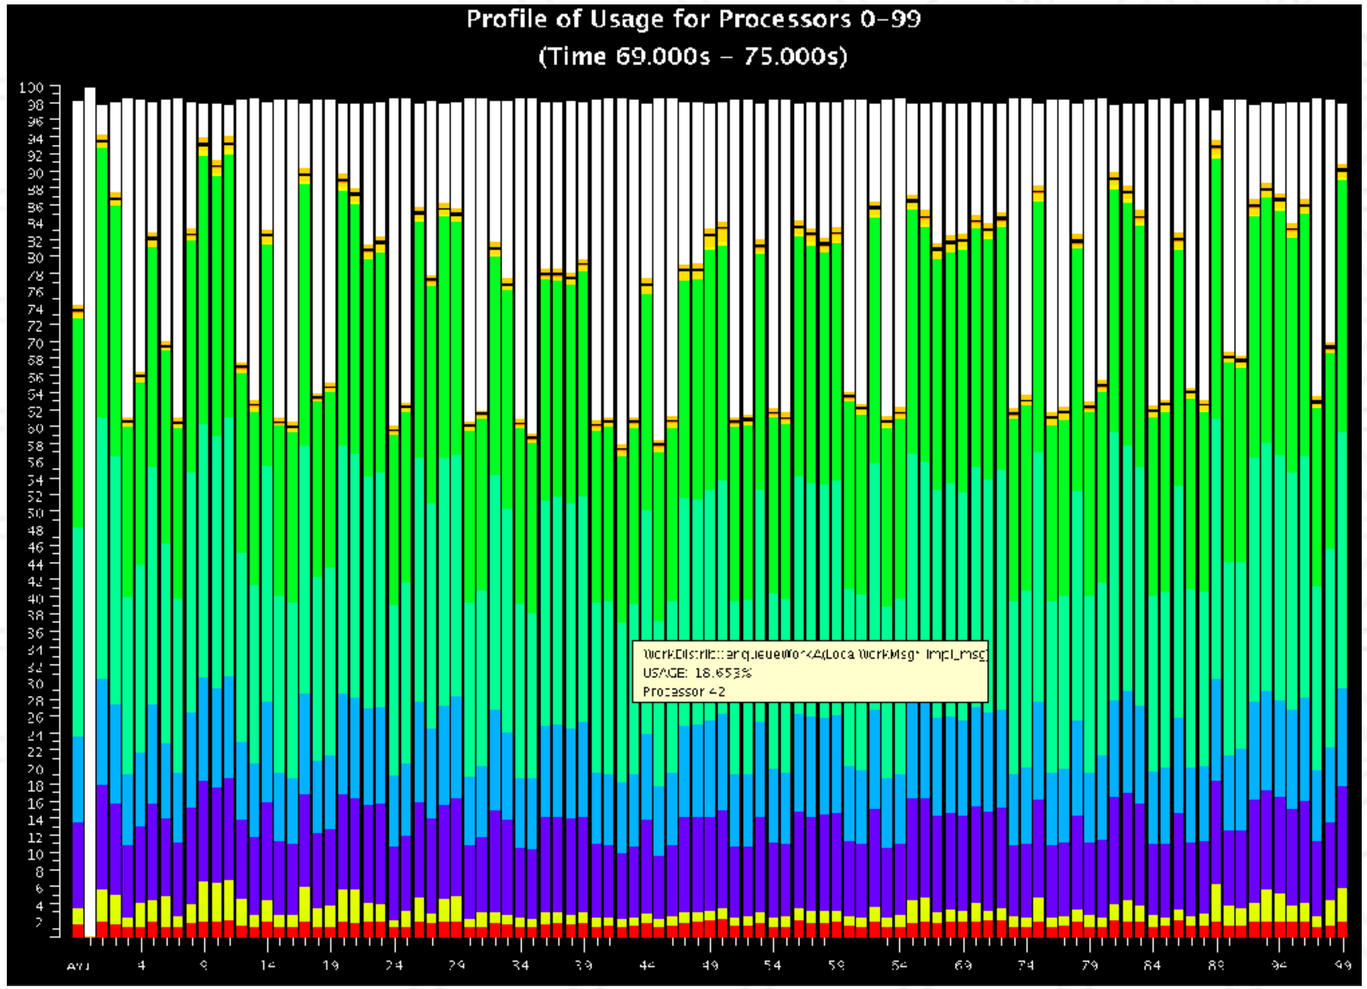

Usage Profile & Histogram View

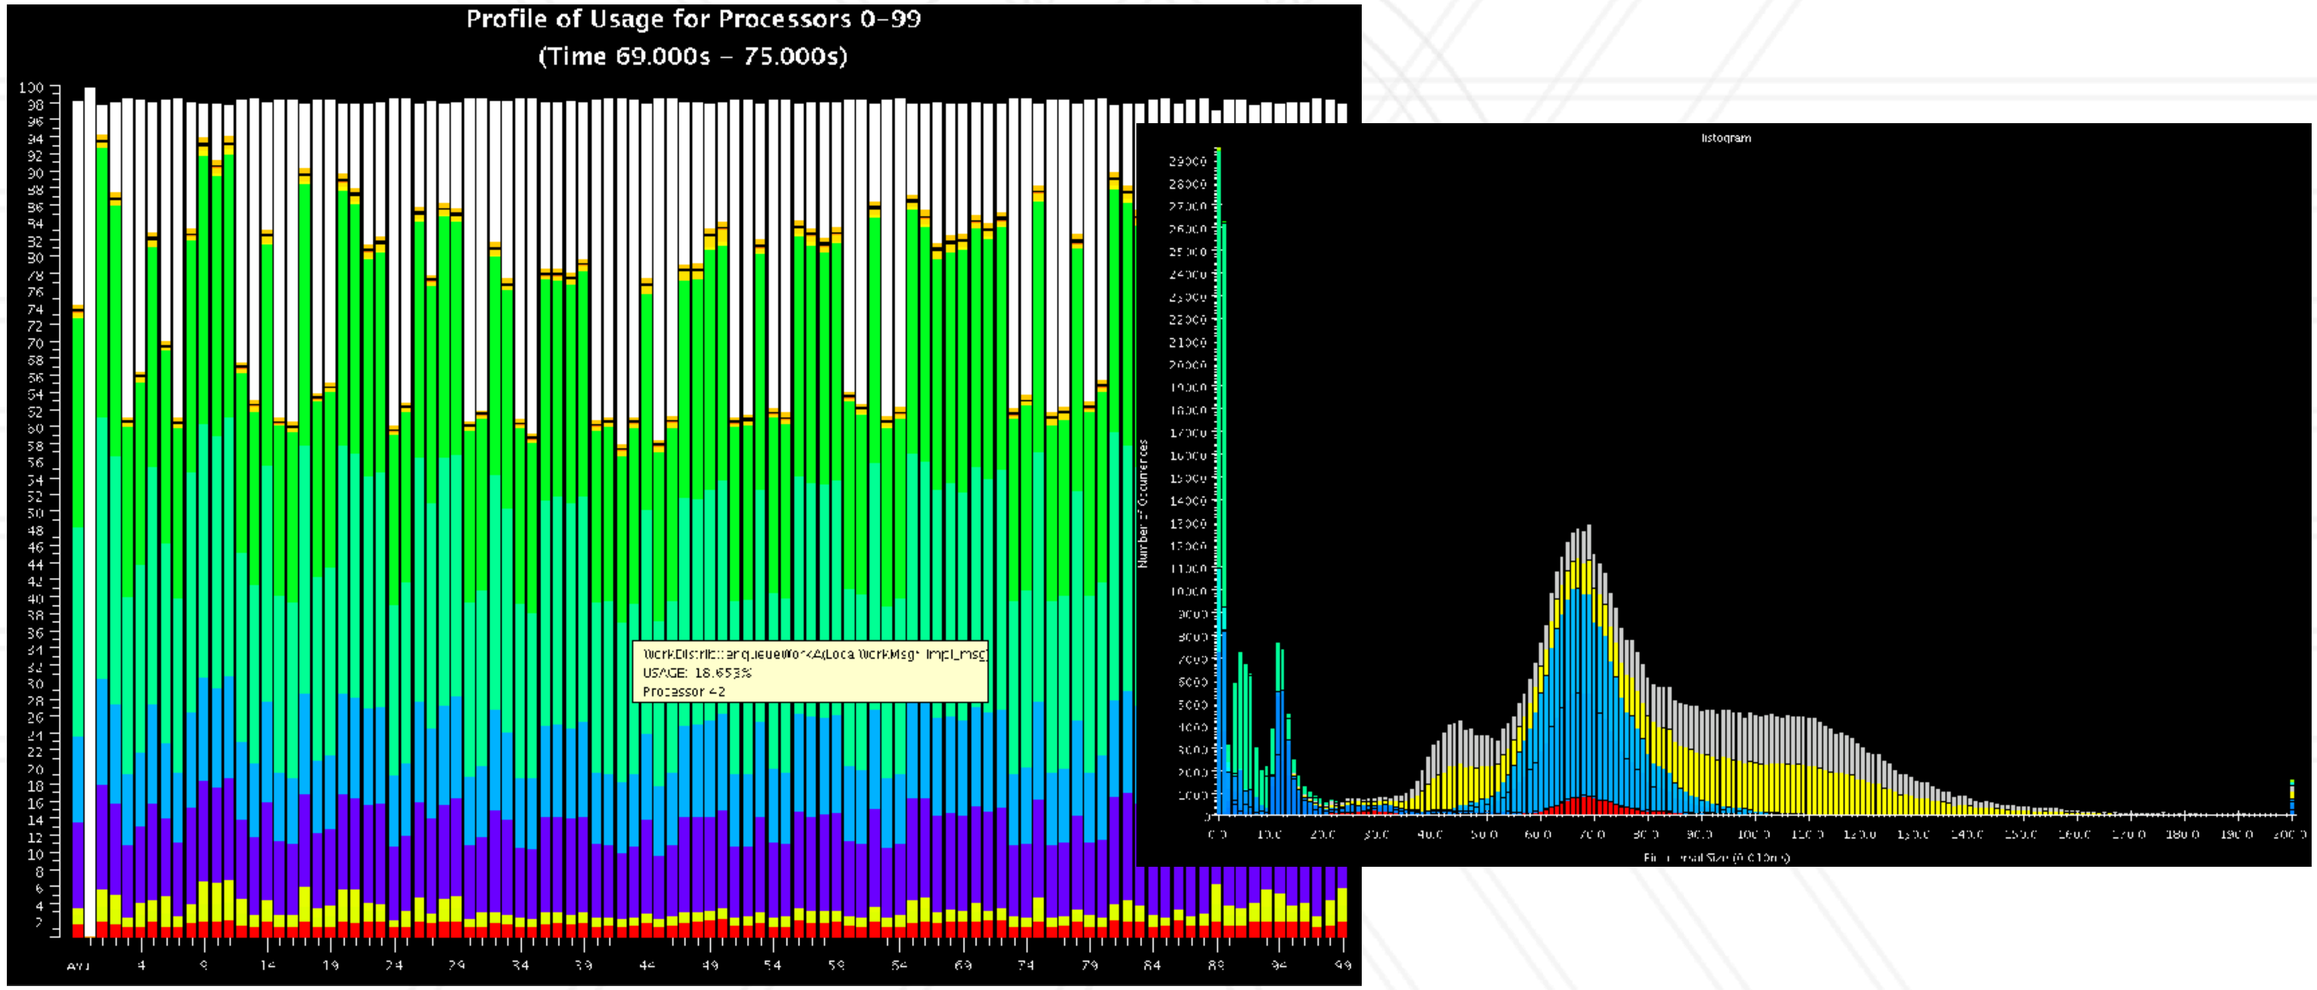

Outlier Analysis

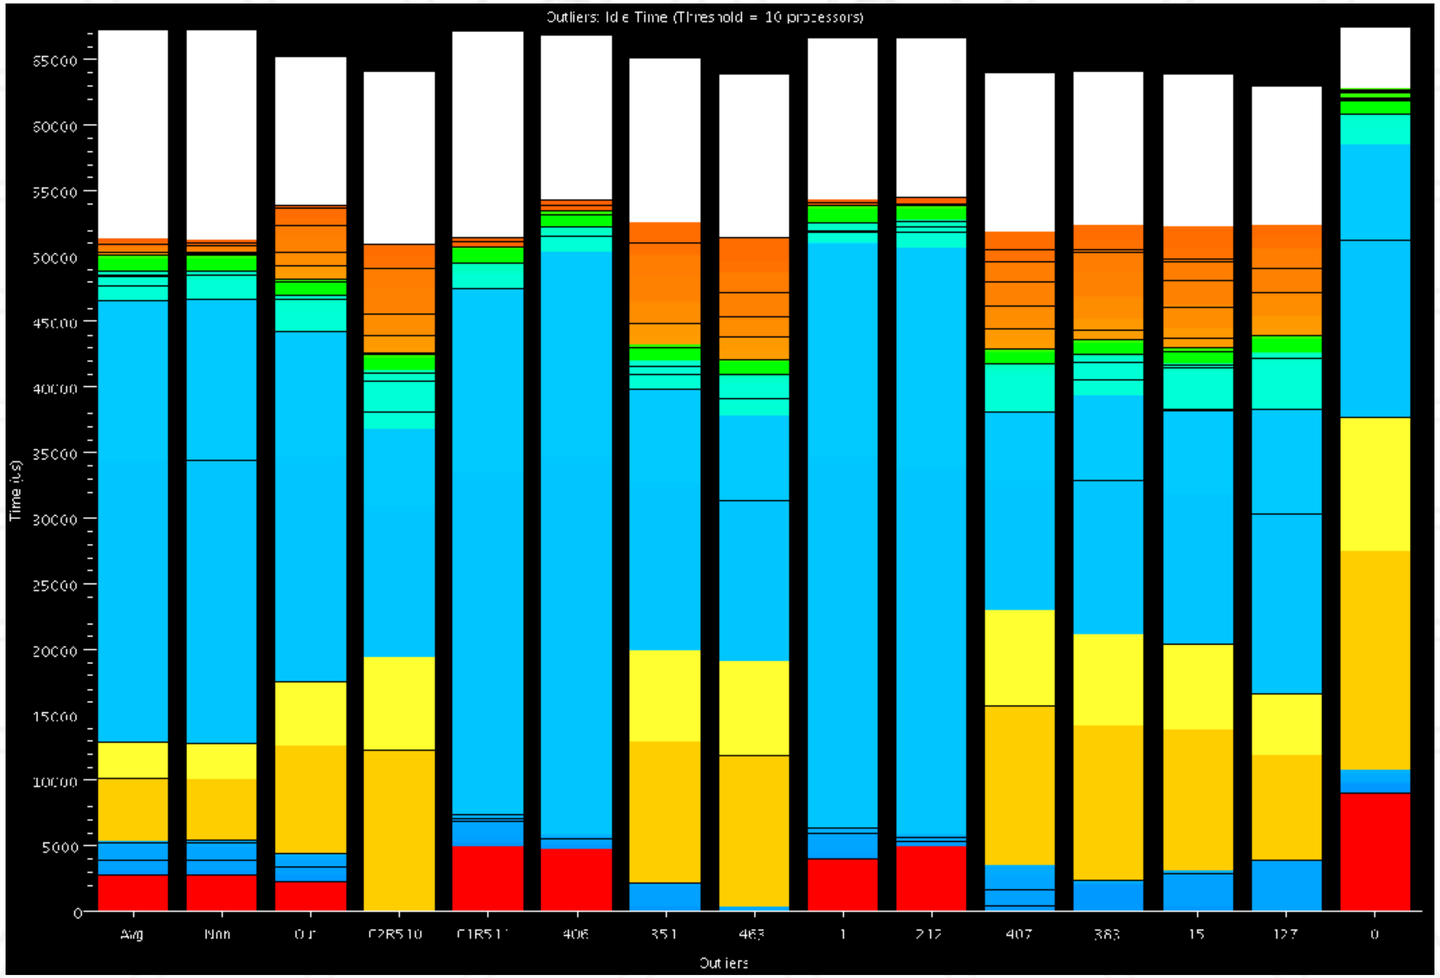

Scripting for multi-run comparisons

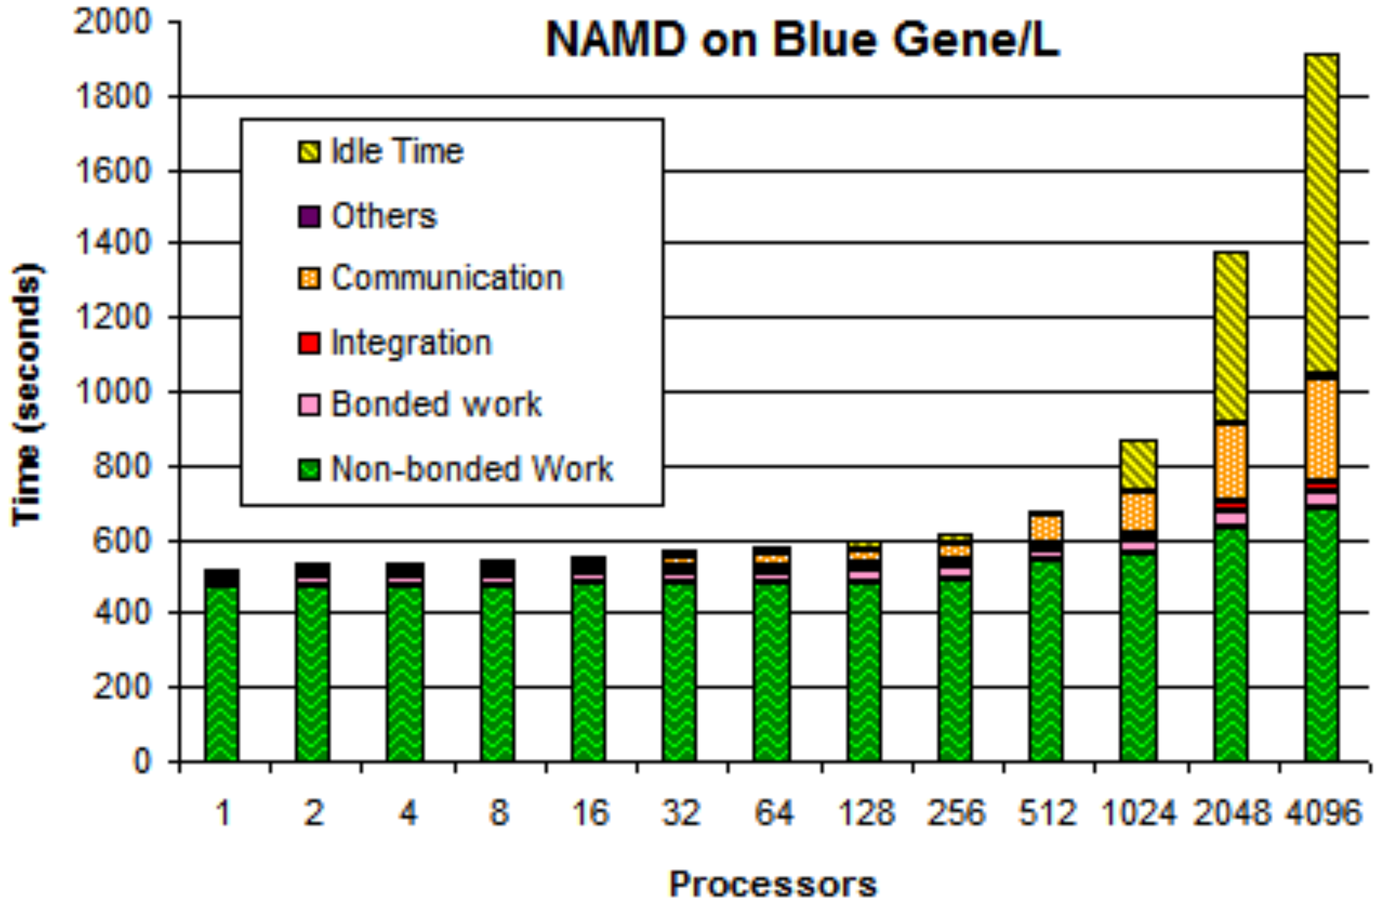

Limitations of current *analysis* tools

- Support their own unique format(s)
- Limited support for saving or automating analysis
- Most tools only support viewing one dataset at a time
- Lack capabilities to sub-select and focus on specific parts

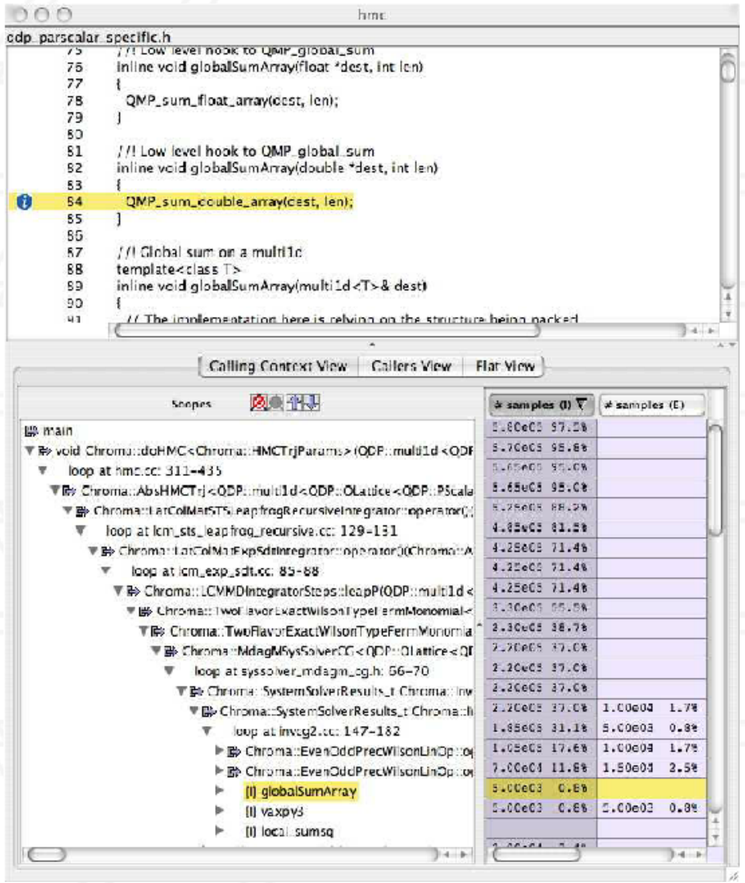

hpcviewer's GUI

Limitations of current *analysis* tools

- Support their own unique format(s)
- Limited support for saving or automating analysis
- Most tools only support viewing one dataset at a time
- Lack capabilities to sub-select and focus on specific parts

Do not enable programmatic analysis of the data by the end user

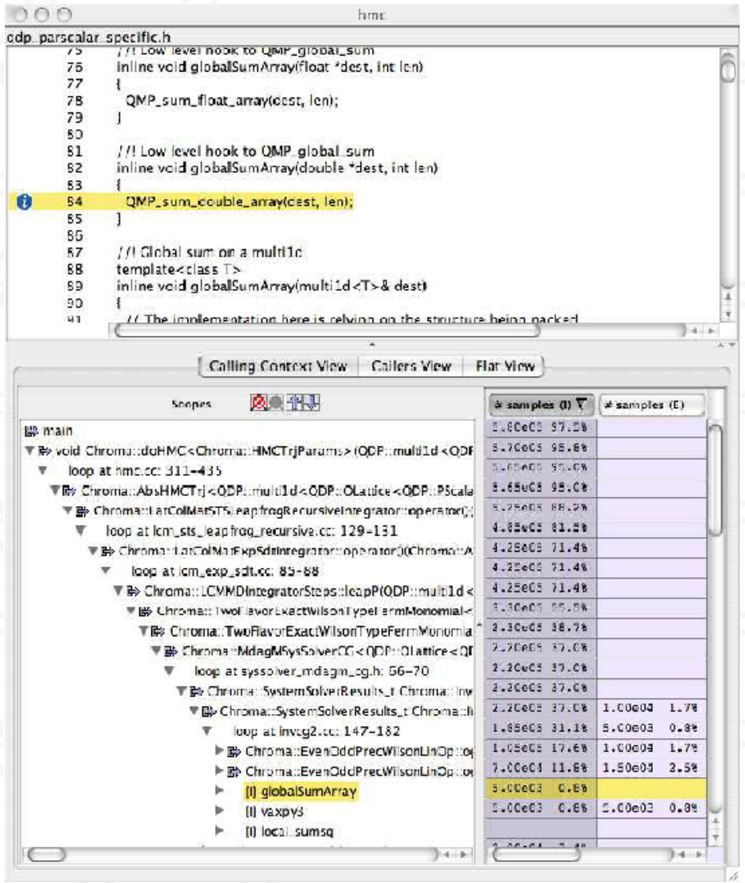

hpcviewer's GUI

Hatchet

- A Python-based library to enable programmatic analysis
- Creates an in-memory representation of the graph
- Leverage pandas which supports multi-dimensional tabular datasets
  - Use graph as structured index to index pandas dataframes
- A set of operators to sub-select and/or aggregate profile data
- A set of operators to compare multiple datasets

Pandas and dataframes

Pandas and dataframes

- Pandas is an open-source Python library for data analysis

Pandas and dataframes

- Pandas is an open-source Python library for data analysis
- Dataframe: two-dimensional tabular data structure
  - Supports many operations borrowed from SQL databases

Columns

Rows

	node	name	time (inc)	time
0	{'name': 'main'}	main	200.0	10.0
1	{'name': 'physics'}	physics	60.0	40.0
2	{'name': 'mpi'}	mpi	20.0	5.0
3	{'name': 'psm2'}	psm2	15.0	30.0
4	{'name': 'solvers'}	solvers	100.0	10.0
5	{'name': 'hypre'}	hypre	65.0	30.0
6	{'name': 'mpi'}	mpi	35.0	20.0
7	{'name': 'psm2'}	psm2	25.0	60.0

Pandas and dataframes

- Pandas is an open-source Python library for data analysis
- Dataframe: two-dimensional tabular data structure
  - Supports many operations borrowed from SQL databases

Index

Columns

Rows

	node	name	time (inc)	time
0	{'name': 'main'}	main	200.0	10.0
1	{'name': 'physics'}	physics	60.0	40.0
2	{'name': 'mpi'}	mpi	20.0	5.0
3	{'name': 'psm2'}	psm2	15.0	30.0
4	{'name': 'solvers'}	solvers	100.0	10.0
5	{'name': 'hypre'}	hypre	65.0	30.0
6	{'name': 'mpi'}	mpi	35.0	20.0
7	{'name': 'psm2'}	psm2	25.0	60.0

Pandas and dataframes

- Pandas is an open-source Python library for data analysis
- Dataframe: two-dimensional tabular data structure
  - Supports many operations borrowed from SQL databases
- MultiIndex enables working with high-dimensional data in a 2D data structure

Index

Columns

Rows

	node	name	time (inc)	time
0	{'name': 'main'}	main	200.0	10.0
1	{'name': 'physics'}	physics	60.0	40.0
2	{'name': 'mpi'}	mpi	20.0	5.0
3	{'name': 'psm2'}	psm2	15.0	30.0
4	{'name': 'solvers'}	solvers	100.0	10.0
5	{'name': 'hypre'}	hypre	65.0	30.0
6	{'name': 'mpi'}	mpi	35.0	20.0
7	{'name': 'psm2'}	psm2	25.0	60.0

Central data structure: a *GraphFrame*

- Consists of a structured index graph object and a pandas dataframe
- Graph stores caller-callee relationships
- Dataframe stores all numerical and categorical data

Central data structure: a *GraphFrame*

- Consists of a structured index graph object and a pandas dataframe
- Graph stores caller-callee relationships
- Dataframe stores all numerical and categorical data

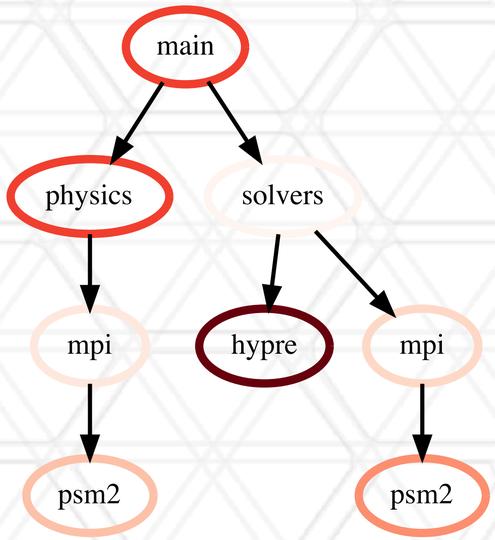

Central data structure: a *GraphFrame*

- Consists of a structured index graph object and a pandas dataframe
- Graph stores caller-callee relationships
- Dataframe stores all numerical and categorical data

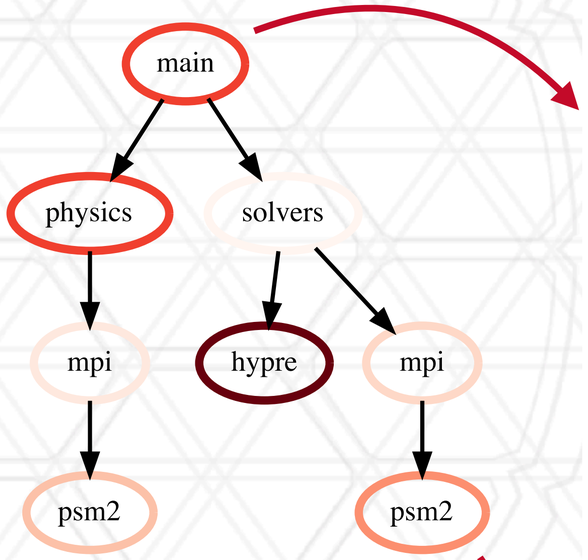

node	name	nid	node	time	time (inc)
main	main	0	main	40.0	200.0
physics	physics	1	physics	40.0	60.0
mpi	mpi	2	mpi	5.0	20.0
psm2	psm2	3	psm2	15.0	15.0
solvers	solvers	4	solvers	0.0	100.0
hypre	hypre	5	hypre	65.0	65.0
mpi	mpi	6	mpi	10.0	35.0
psm2	psm2	7	psm2	25.0	25.0

Questions?

UNIVERSITY OF
MARYLAND

Abhinav Bhatele

5218 Brendan Iribe Center (IRB) / College Park, MD 20742

phone: 301.405.4507 / e-mail: bhatele@cs.umd.edu