Hatchet: Pruning the Overgrowth in Parallel Profiles

Abhinav Bhatele^{†,*}, Stephanie Brink^{*}, and Todd Gamblin^{*}

[†]Department of Computer Science, University of Maryland, College Park

^{*}Center for Applied Scientific Computing, Lawrence Livermore National Laboratory

bhatele@cs.umd.edu, brink2@llnl.gov, tgamblin@llnl.gov

Figure 1: A flame graph representation of a user-annotated nested region profile of the LULESH proxy application.

ABSTRACT

Performance analysis is critical for eliminating scalability bottlenecks in parallel codes. There are many profiling tools that can instrument codes and gather performance data. However, analytics and visualization tools that are general, easy to use, and programmable are limited. In this paper, we focus on the analytics of structured profiling data, such as that obtained from calling context trees or nested region timers in code. We present a set of techniques and operations that build on the pandas data analysis library to enable analysis of parallel profiles. We have implemented these techniques in a Python-based library called Hatchet that allows structured data to be filtered, aggregated, and pruned. Using performance datasets obtained from profiling parallel codes, we demonstrate performing common performance analysis tasks reproducibly with a few lines of Hatchet code. Hatchet brings the power of modern data science tools to bear on performance analysis.

CCS CONCEPTS

• **General and reference** → Performance; • **Software and its engineering** → Software maintenance tools.

KEYWORDS

performance analysis, tool, parallel profile, calling context tree, call graph, graph analytics

ACM Reference Format:

Abhinav Bhatele, Stephanie Brink, and Todd Gamblin. 2019. Hatchet: Pruning the Overgrowth in Parallel Profiles. In *The International Conference for*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SC '19, November 17–22, 2019, Denver, CO, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6229-0/19/11...\$15.00

<https://doi.org/10.1145/3295500.3356219>

High Performance Computing, Networking, Storage, and Analysis (SC '19), November 17–22, 2019, Denver, CO, USA. ACM, New York, NY, USA, 11 pages.
<https://doi.org/10.1145/3295500.3356219>

1 MOTIVATION

Understanding performance bottlenecks is critical to optimizing the performance of high performance computing (HPC) codes. Profiling tools [3, 5, 12, 21, 25] allow developers to focus their optimization efforts by pinpointing the parts within a code's execution that consume the most time. Without them, it would be extremely difficult to find performance problems, especially in modern applications, which can comprise millions of lines of code.

Attributing time to code can be tricky, and it requires a reasonable understanding of the structure of the program. Most basic profilers attribute time to functions or statements in the code. More sophisticated profilers can record time spent in different *call paths* or *calling contexts*. For example, the profiler would differentiate time spent in MPI_Send when it is called in a hydrodynamics routine from time spent in MPI_Send when it is called in a solver library. Such profilers maintain a prefix tree of unique calling contexts. Other profilers may attribute time to nodes in a static or dynamic *call graph*, in which case time would be attributed to nodes in the graph. So, code regions can be represented by a range of structures, from simple, “flat” strings, to nodes in trees or graphs. Figure 1 shows an example of a simple tree, where each node (each box in the flame graph) represents a code region annotated by the user.

Most profiling tools use their own unique format to store recorded data, and they may display this data as text or with a tool-specific viewer (typically a GUI). These tools are limited in the kinds of analysis they support, and they do not enable the end user to *programmatically* analyze performance data. Most profile data viewers provide a point-and-click-based workflow, with limited support for saving or automating analysis. As such, analyzing performance data can be very tedious, and using a new *measurement* tool often requires also using a new *analysis* tool to understand the data. Moreover, most tools only support viewing one or two call graphs

at a time. They are often insufficient for tasks like load balance analysis that require detailed averaging and clustering across ranks, threads, and time. They also lack general capabilities for effectively sub-selecting and focusing on specific *parts* of larger datasets.

Growing interest in data science has led to the availability of many fully-featured data analysis environments. Python and R in particular support *DataFrames* that blend features of traditional numerical computing environments such as MATLAB with database-derived features such as indexed queries, joins, and aggregations. Numerous plotting libraries and visualization tools are available to view data in these environments, all with programmable APIs. In addition, data analysis tools have more features, and are better maintained by open-source software communities than tool GUIs. While these environments can handle numerical and temporal indices with ease, they cannot handle structured datasets, such as profiles that are indexed by nodes in a tree or a graph.

This paper presents a set of techniques that allow modern data analysis libraries to be leveraged for parallel profile analysis. We have developed a canonical data model that can read, represent, and index the data generated by most profiling tools. We call this data model a *structured index*. We have developed techniques for selecting, filtering, and aggregating datasets with structured indices, and we have generalized these techniques for datasets that have hybrid indices over code, processes, threads, and time. Finally, we have implemented these techniques in a library that we call *Hatchet*, which builds on the popular pandas data analysis library [15, 16].

This paper makes the following contributions:

- A canonical data model that enables different types of profile data (e.g., HPCToolkit, Caliper, callgrind, gprof, perf) to be represented and analyzed in a common way;
- an indexing technique that allows structured graph or tree nodes to be used and queried as a DataFrame index;
- operators that allow structured data to be filtered, aggregated, and pruned to produce API-centric sub-profiles; and
- an implementation of these techniques in the Hatchet library that enables users to leverage modern data analysis approaches for analyzing large-scale call path profiling data.

Using performance datasets derived from running HPC codes, we also present several case studies that demonstrate performing common performance analysis tasks reproducibly with only a few lines of Hatchet code. Examples include: 1) identifying regions or callsites with the most load imbalance across MPI processes or threads; 2) filtering datasets by a metric or library/function names to focus on subsets of data; and 3) easily handling and analyzing multi-rank, profile data from multiple executions. We expect that Hatchet will make analysis of HPC performance data quicker, easier, and more reproducible.

2 BACKGROUND

We define the different kinds of structured profiles output by various profiling tools (in particular HPCToolkit and Caliper), and provide a brief background on the pandas data analysis library.

2.1 Structured Performance Data

Profiles can be collected either through *sampling* or through direct instrumentation. In a directly instrumented program, measurement

code is typically inserted with a compilation tool, and data is collected at each instrumentation point. Sampled profilers instead periodically force an interrupt while a program runs, and data is collected at each interrupt. In either scenario, the collected data contains two types of information: *contextual information*, i.e., the current line of code, file name, the call path, the process ID, etc.; and *performance metrics*, such as the number of floating point operations or branch misses that occurred since the last sample. Depending on how these samples are aggregated, different types of profiles can be generated.

Calling Context Trees (CCT): *Callpath profilers* analyze the stack at runtime to determine the full *calling context* of each sample. Calling contexts, or call paths, refer to the sequence of function invocations that led to the sampled one. Each stack frame becomes a node in the CCT, and the path from the root of the tree to a given node represents the call path, or calling context, that led to the leaf invocation. A *calling context tree (CCT)* is a prefix tree of call paths. CCTs allow analysts to understand differences in function behavior that depend on *how* the function was called.

Call Graphs: *Call graph profiles* [7] do not perform the stack analysis required to generate CCTs; they attribute data only to the *name* of each function called. Edges in the call graph represent *static* calling relationships (i.e., that one function is called by another), averaged across all invocations regardless of origin. Samples can also be aggregated at the granularity of user-annotated regions, which can produce an even coarser tree or graph than call graphs. Call graphs are a more concise, (and sometimes clearer) representation than a CCT but they discard all context information.

The most insightful representation of a given performance profile depends on the problem being analyzed. Hatchet's data model is designed to handle structured profiling data generated at various granularities from different file formats. To motivate this, we provide a brief overview of two profiling tools in the next section.

2.2 Call Path Profilers

HPCToolkit: HPCToolkit [3] provides a suite of performance tools enabling measurement, analysis, correlation, and visualization. When asynchronous or synchronous events occur in the application, HPCToolkit records the full calling context as a CCT. With this data structure, a unique call path for a given node corresponds to the path from that node to the root. In HPCToolkit, a CCT node is not limited to function invocations only, but can also record loops, statements, and other code structures. Moreover, for parallel programs, HPCToolkit records the metrics per process and thread for every node in a unified CCT. The database generated by HPCToolkit's `hpcprof-mpi` utility is derived by modifying the CCT to include the node ID and timestamp. Since there can be multiple processes in a parallel run, `hpcprof-mpi` outputs the unified tree in XML format as well as separate binary files for each process that contain the metrics for all the nodes in the CCT.

Caliper: Caliper [5] provides a general abstraction layer for application performance introspection. Application developers can use Caliper's annotation APIs to collect performance information. It has a flexible data aggregation model [4] designed to handle a

wide variety of information that can be analyzed offline or in situ. At runtime, Caliper builds a generalized context tree consisting of attributes representing data elements. The context information for any given node can be derived by collecting all attributes on the path between the node and the root node. User annotations in the code or enabling the call path service in Caliper can generate a structured profile or CCT. Caliper supports JSON output formats that can be generated by either running `cali-query` on the raw Caliper samples or by enabling the `mpireport` service.

2.3 Pandas and DataFrames

Pandas: *pandas* is an open-source Python library providing data structures and tools for data analysis [15, 16]. It is built on top of NumPy and is well-suited for various kinds of data including tabular and statistical datasets. Pandas provides two main data structures: *Series* and *DataFrame*.

Series: A Series is a one-dimensional, homogeneously-typed array that has an associated *index*. Unlike a traditional array, the index need not be numerical; a Series can be indexed by any ordered, comparable data type. In this sense a Series in pandas is somewhat like a sorted dictionary or hashtable, as it allows fast lookup by non-numerical values.

DataFrame: A DataFrame is a two-dimensional tabular structure with potentially heterogeneously-typed columns. Each column in a DataFrame can be thought of as its own Series, and certain columns can be made the *index* of the DataFrame. Columns have titles, or *labels* that can be used to retrieve their corresponding Series, and the values in index columns can be thought of as keys for fast lookup of rows in the DataFrame. DataFrames with like indices can thus be aligned and compared. In addition to indexing, DataFrames support a wide range of functionality borrowed from the world of spreadsheets and SQL databases: operations to handle missing data, slicing, fancy indexing, subsetting, inserting and deleting columns, merging datasets, aggregating or grouping data, etc. Most importantly, multiply indexing a DataFrame provides an intuitive way of working with high-dimensional data in a two-dimensional data structure.

MultiIndex: pandas allows designating multiple columns in the a DataFrame to be a composite *MultiIndex*. This enables users to easily store and manipulate data with an arbitrary number of dimensions. For example, in parallel performance analysis, we may want to index data not only by a calling context or other structured code identifier, but also by MPI rank, node hostname, or thread id. Pandas makes this natural; we can easily add data from additional levels of parallelism by adding additional index columns to a MultiIndex.

3 THE HATCHET LIBRARY

We have created Hatchet, a Python library that builds on top of pandas so that it can deal with structured profiling data such as calling context trees (CCTs) and call graphs. As mentioned in the previous section, pandas provides easy ways to manipulate data in Series and DataFrames, and it has support for arbitrary-dimensional indexes through MultiIndexes. However, pandas cannot handle structured datasets such as profiles that refer to source code, and

are indexed by nodes in a tree or a graph. Pandas by default handles indexes of numbers, text, or dates, but these are all essentially linear data spaces. Call trees and call graphs have nonlinear node and edge structures, and we want to be able to preserve the ability to reason about graph- and tree-based relationships like parent, ancestor, and child. To overcome this, Hatchet provides data structures that enable indexing rows of a DataFrame by nodes in the graph.

3.1 Structured Indexes

We have developed a canonical data model that can represent and *index* the data generated by most profiling tools. We call this data model a *structured index*. This index enables nodes in a structured graph or tree to be used as a DataFrame index. There can be multiple types of nodes, such as procedure/function nodes, loop nodes (representing loop structures), and statement nodes (leaf-level nodes).

Hatchet’s structured index is, at the most basic level, an in-memory graph. The structure of the graph is shown in Figure 2. This particular graph happens to be a tree, but Hatchet can support both call graphs and call trees. In the example graph, a function A calls a function B and a function C, and B contains a loop nest, D. These code structures span two libraries, or “modules”, `libfoo.so` and `libbar.so`.

Figure 2: Hatchet’s graph object, showing nodes and Frames.

3.1.1 Generic Frames. Each node in Hatchet’s graph contains a generic identifier for the code construct it represents. We call this identifier a *Frame* (after a *call frame*). Frames are generated by data sources such as file readers; they contain a set of key/value pairs that describe the source code the node represents. Depending on the type and amount of data the reader provides, this may be as simple as a function name or region name. It may also be more complex, including file and line number, code module, or other data from the particular input tool. Frames support basic comparison operations, such as `==`, `>`, `<`, etc., which are evaluated based on the names and values of component fields. If two Frames do not have completely identical keys *and* values, they are not considered equal. The key design point here is that there is not a rigid schema for the Frames; they are *generic* and can be generated by each data source according to its granularity. Hatchet attempts to regularize

the *names* of fields as much as possible over different sources to enable comparison of data across measurement tools, but this is not a requirement.

3.1.2 Nodes. Frames are associated with nodes in the Hatchet graph, and node objects define connectivity and structure of the Hatchet model. Each node knows its children and its ancestors in the graph, and each node has a unique *key*. The key is not meant to be accessed by Hatchet users. Rather, like Frames, Hatchet nodes expose their *own* comparison operations (`==`, `>`, `<`, etc.), which opaquely operate on this key. This means that we can insert Node objects directly into a pandas DataFrame column and make it an *index*. By default, we use the Python `id()` function for the node key. This is equivalent, roughly, to C's `&` operator, in that it returns an integer representing the address of the Python object in memory. We require only that the node key be *unique* for each node. We can optionally use keys that provide certain useful orderings (like pre-order, post-order, etc.), if we want to pay the cost of a graph traversal (or sort) to generate more structured keys. We default to only guaranteeing uniqueness and not order in our keys.

3.2 GraphFrame

The central data structure in the Hatchet library is a *GraphFrame*, which combines the structured index *Graph* with a pandas *DataFrame*. Figure 3 shows the two objects in a GraphFrame – a graph object (the index), and a DataFrame object storing the metrics associated with each node.

Figure 3: In Hatchet, the GraphFrame consists of a graph and a DataFrame object.

Because of the way we have architected the structured index Graph, we can insert Node objects directly into the pandas DataFrame. The nodes are sorted using their basic comparison operators, which operate on their key attribute. Thus, the first column in the DataFrame (the node) is the index column. As a convenience, we may also add columns (like *name*) based on attributes from each node's Frame. For example, in the figure, we have added the *name* and *nid* columns from the Frame subclass for HPCToolkit. This allows us to use regular pandas operations (selection, filtering) on these values directly. As we will see later, the *node* column itself also allows various graph-semantic functions to be used, as well. Finally, in addition to the identifying information for each node, we *also* add columns for each associated performance metric (inclusive and exclusive time in the figure).

Graphs vs. Trees: Hatchet stores the structure (typically a prefix tree generated from call paths) in the input data as a directed graph (instead of a tree) for two reasons. First, subsequent operations on a tree can create edges or merge nodes, turning the tree into a graph. Additionally, output from tools such as callgrind is already in the form of a DAG. Hatchet's directed graph could be connected or have multiple disconnected components. Each entity in the graph, such as a callsite, procedure frame, or function, is stored as a node and the caller-callee relationships are stored as directed edges. Each node in the graph can have one or multiple parents and children.

Benefits of DataFrames: We use a pandas DataFrame to store all the numerical and categorical data associated with each node. Profile data can be inherently high-dimensional when metrics are gathered per-MPI process and/or per-thread. In such cases, each node in the call tree or graph has metrics per-MPI process and/or thread and this data needs to be stored and indexed hierarchically. To index the rows of the data frame in such cases, a MultiIndex consisting of the structured index for the node and MPI rank or thread ID is used. In the most general case, a row in the data frame is indexed by a process and/or thread ID (and any other needed identifiers in even higher dimensional cases).

3.3 Immutable Graph Semantics

Astute readers may have noted that we are adding direct references to graph nodes into the DataFrame. The risk this poses in our API is that client code can extract a subset of a DataFrame and hand it off to *other* client code, which then modifies the graph index nodes directly and corrupts all DataFrames that use the same nodes. One key aspect of Hatchet is that its graph nodes use *immutable* semantics. The GraphFrame API is responsible for ensuring that operations *between* any two GraphFrames use immutable graph node references, and that any operations that would modify a graph node in place instead create an *entirely new* graph index for the new GraphFrame to work with. So, we implement immutable semantics using copy-on-write to simplify the management of the graph and DataFrame together.

One further consequence of our index model is that to use two DataFrames *together*, we require that their graphs be *unified*. That is, that they share the same index. This should be obvious when considering that the nodes are compared by their key values, and two nodes can only be considered identical within an index if they have identical keys, which means that they *must* be in the same graph for comparison to make sense. We accomplish this by traversing the graphs and computing their *union* according to their connectivity and Frame values (described further in the API section). Incidentally, this type of restriction is not unusual in pandas, where comparing two data frames frequently requires reconciling their indexes, as well. We abstract the details of these graph operations in Hatchet through the GraphFrame API, which determines when and how GraphFrames should be unified.

3.4 Reading a CCT Dataset

With all of these components, the structured index Graph models the edge relationships between nodes in the structured data, *and* a DataFrame stores the numerical (performance metrics such as time, performance counter data, etc.) and categorical data (e.g., load

module, file, line number) associated with each node. The generality of what can be stored in a pandas DataFrame enables Hatchet to store almost any kind of contextual information recorded during sampling by diverse profiling tools.

Hatchet provides readers for several input formats to support data collected by popular profiling tools in the HPC community. Hatchet can read in the database directory generated by HPCToolkit (hpcprof-mpi), and also JSON files generated by Caliper. In addition, one can provide structured data in the Graphviz DOT format, or simple dictionary and list literals in Python.

Most profiling tools that generate CCTs have two kinds of information in their output, often separated into different parts of a file or different files. The first information is the structure of the CCT – present in `experiment.xml` in HPCToolkit databases, and the nodes section of a Caliper JSON file. The second piece of information is the performance metrics attached to each node – available in `metric-db` files in HPCToolkit data and in the `data` section of a Caliper JSON file. The readers in Hatchet read in both pieces of information. The CCT structure is used to construct the graph object of the GraphFrame and the performance metrics are used to construct the DataFrame object. As the readers construct these two objects, they also make connects between the graph and DataFrame objects using the structured index.

4 THE HATCHET API

We now describe some of the operators provided by the Hatchet API that allow structured data to be manipulated in different ways: filtered, aggregated, pruned, etc. Even though all of the operations below are performed on the GraphFrame, some only modify the DataFrame, some only modify the graph, and others modify both. They are categorized accordingly in the following sections. Note that we consider a graph to be immutable, so any operations that lead to changes in the graph structure will create a new graph and return a new GraphFrame indexed by the new graph's nodes.

4.1 DataFrame Operations

filter: Filter takes a user-supplied function and applies that to all rows in the DataFrame. The resulting Series or DataFrame is used to filter the DataFrame to only return rows that are true. The returned GraphFrame preserves the original graph provided as input to the filter operation. Figure 4 shows a DataFrame before and after a filter operation. In this case, the applied function returns all rows where time is greater than 10.0.

Filter is one of the operations that leads to the graph object and DataFrame object becoming inconsistent. After a filter operation, there are nodes in the graph that do not return any rows when used to index into the DataFrame. Typically, the user will perform a squash on the GraphFrame after a filter operation to make the graph and DataFrame objects consistent again.

drop_index_levels: When there is per-MPI process or per-thread data in the DataFrame, a user might be interested in aggregating the data in some fashion to analyze the graph at a coarser granularity. This function allows the user to drop the additional index columns in the hierarchical index by specifying an aggregation

	name	nid	node	time	time (inc)
node					
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
hypr	hypr	5	hypr	65.0	65.0
mpi	mpi	6	mpi	10.0	35.0
psm2	psm2	7	psm2	25.0	25.0

	name	nid	node	time	time (inc)
node					
main	main	0	main	40.0	200.0
physics	physics	1	physics	40.0	60.0
psm2	psm2	3	psm2	15.0	15.0
hypr	hypr	5	hypr	65.0	65.0
psm2	psm2	7	psm2	25.0	25.0


```

1 gf = GraphFrame( ... )
2 filtered_gf = gf.filter(lambda x: x['time'] > 10.0)

```

Figure 4: The DataFrame before (left) and after (right) a filter operation on the time column.

Figure 5: A graph before (left) and after two squashes (middle, right) on the GraphFrame.

function. Essentially, this performs a groupby and aggregate operation on the DataFrame. The user-supplied function is used to perform the aggregation over all MPI processes or threads at the per-node granularity.

update_inclusive_columns: When a graph is rewired (i.e., the parent-child connections are modified), all the columns in the DataFrame that store inclusive values of a metric become inaccurate. This function performs a post-order traversal of the graph to update all columns that store inclusive metrics in the DataFrame for each node.

4.2 Graph Operations

squash: The squash operation is typically performed by the user after a filter operation on the DataFrame. As shown in Figure 5, the squash on line 2 removes nodes from the graph that were previously removed from the DataFrame due to a filter operation. When one or more nodes on a path are removed from the graph, the nearest remaining ancestor is connected by an edge to the nearest remaining child on the path. All call paths in the graph are re-wired in this manner.

In some cases, a squash may need to merge nodes. The filter and squash calls on lines 4-6 remove the physics and hypr nodes from the graph, but main must now connect to *both* psm2 nodes. In a calling context tree, a node cannot have two children with identical frames, so we merge the psm2 nodes together. The graph now represents *only* the time spent in psm2 when called directly *or* transitively from main. As mentioned earlier, node merging can convert a tree into a graph. Squash and other Hatchet API calls are general and handle both trees and graphs.

A squash operation creates a new DataFrame in addition to the new graph. The new DataFrame contains all rows from the original DataFrame, but its index points to nodes in the new graph. Additionally, a squash operation will make the values in all columns containing inclusive metrics inaccurate, since the parent-child relationships have changed. Hence, the squash operation also calls `update_inclusive_columns` to make all inclusive columns in the DataFrame accurate again.

equal: This checks whether two graphs have the same nodes and edge connectivity when traversing from their roots. If they are equivalent, it returns true, otherwise it returns false.

union: The union function takes two graphs and creates a unified graph, preserving all edges structure of the original graphs, and merging nodes with identical context. When Hatchet performs binary operations on two GraphFrames with unequal graphs, a union is performed beforehand to ensure that the graphs are structurally equivalent. This ensures that operands to element-wise operations like add and subtract, can be aligned by their respective nodes.

4.3 GraphFrame Operations

copy: The copy operation returns a shallow copy of a GraphFrame. It creates a new GraphFrame with a copy of the original GraphFrame's DataFrame, but the *same* graph. As mentioned earlier, graphs in Hatchet use immutable semantics, and they are copied only when they need to be restructured. This property allows us to reuse graphs from GraphFrame to GraphFrame if the operations performed on the GraphFrame do not mutate the graph.

unify: Similar to union on graphs, unify operates on GraphFrames. It calls union on the two graphs, and then reindexes the DataFrames in both GraphFrames to be indexed by the nodes in the unified graph. Binary operations on GraphFrames call unify which in turn calls union on the respective graphs.

add: Assuming the graphs in two GraphFrames are equal, the add (+) operation computes the element-wise sum of two DataFrames. In the case where the two graphs are not identical, `unify` (described above) is applied first to create a unified graph before performing the sum. The DataFrames are copied and reindexed by the combined graph, and the add operation returns new GraphFrame with the result of adding these DataFrames. Hatchet also provides an in-place version of the add operator: `+=`.

subtract: The subtract operation is similar to the add operation in that it requires the two graphs to be identical. It applies union and reindexes DataFrames if necessary. Once the graphs are unified, the

```
1 gf1 = GraphFrame( ... )
2 gf2 = GraphFrame( ... )
3
4 gf2 -= gf1
```

Figure 6: Subtraction operation on two GraphFrames (resulting graph at the bottom).

subtract operation computes the element-wise difference between the two DataFrames. The subtract operation returns a new GraphFrame, or it modifies one of the GraphFrames in place in the case of the in-place subtraction (`-=`). Figure 6 shows the subtraction of one GraphFrame from another and the graph for the resulting GraphFrame.

4.4 Visualizing Output

Hatchet provides its own visualization as well as support for two other visualizations of the structured data stored in the graph object. The native visualization in Hatchet is a string that can be printed to the terminal to display the graph. Hatchet can also output the graph in the DOT format or a folded stack used by flame graph [8].

The dot utility in Graphviz produces a hierarchical drawing of directed graphs, particularly useful for showing the direction of the edges. Flame graphs are useful for quickly identifying the performance bottleneck, that is the box with the largest width. The y-axis of the flame graph represents the call stack depth. Figure 7 shows the same Hatchet graph presented in the three supported visualizations: terminal, DOT, and flame graph. For particularly large graphs, these visual representations can be useful for quickly identifying caller-callee relationships. However, identifying performance bottlenecks or load imbalance might be easier in the DataFrame.

5 PERFORMANCE

It is vital that performance analysis tools have low overheads and that they enable quick analysis of performance datasets without the user having to wait for a long time for each operation to complete. In Figure 8, we provide execution times for some operations in

Figure 7: Visualization outputs supported in Hatchet include terminal output (left), DOT (right), and flame graph (bottom).

Hatchet when using increasingly large datasets. We ran LULESH to generate Caliper profiles on 1 to 512 cores. LULESH requires a cubed number of processes. Hatchet was run on a relatively slow macOS laptop (1.8 GHz Intel Core i5). In the plot, *file read* is the time to read the input dataset into memory and convert it into the Hatchet data representation (graph and DataFrame). *drop index* represents the `drop_index_levels` operation, which we use to aggregate the per process information. If we apply a filter after dropping the second index (MPI rank), the filter operation takes a constant amount of time (~ 0.2 seconds). Hence, in the plot, the time shown for filter is measured for the case when filter is done without aggregating the per-process information. We see that the time increases linearly with the increase in the size of the dataset (both axes have a logarithmic scale).

Hatchet only adds a modest amount of code on top of the pandas library. Currently, the Hatchet code is nearly 2,400 lines of Python (obtained using `sloccount` [26]). We expect it to grow modestly as we add more readers and operations to it.

6 CASE STUDIES

In this section, we present several case studies demonstrating how common performance analyses can be executed in an automated manner using the Hatchet API and a few lines of Python code. The first set of case studies analyze single execution profiles for two scientific proxy applications, while the second set of case studies compare profiles from multiple executions.

6.1 Experimental Setup

We performed our single- and multi-node experiments on the Quartz supercomputer at Lawrence Livermore National Laboratory (LLNL). Each node of Quartz contains two Intel Broadwell processors with 36 cores per node. Our case studies used two scientific

Figure 8: Performance overheads for different operations in Hatchet shown on a logarithmic scale. *file read* is the time to convert the data into the Hatchet representation, *drop index* and *filter* are the time to complete the `drop_index_levels` and `filter` operations, respectively.

proxy applications. LULESH [1] is a Lagrangian shock hydrodynamics mini-application that solves a Sedov blast problem. For these case studies, we instrumented the LULESH code with Caliper annotations to collect performance metrics in Caliper’s split JSON format. The second proxy application we used was Kripke [2, 13], which simulates neutron transport. We used HPCToolkit to generate the execution profiles of Kripke.

6.2 Analyzing a Single Execution Profile

Analyzing the profiling output from a single application execution is a fairly common performance analysis task. Typically, end users or performance researchers profile their code on a platform using a number of processes where they expect or have witnessed a performance degradation, and then analyze the output of such profiling. One of the most common tasks is to pin-point the regions of code or functions where the code spends most of its time. This is traditionally called a flat profile because the calling context is lost and we just get a flat view of functions or statements or code regions.

Flat profiles: Flat profiles can be easily generated in Hatchet using the groupby functionality in pandas. The flat profile can be based on any categorical column (e.g., function name, load module, file name). Similar to the sort feature in `perf`, the flat profile groups the nodes by the specified categorical column. Figure 9 shows the code to generate a flat profile by applying a groupby operation on the DataFrame object. The data read into Hatchet was generated by profiling 20 time steps of Kripke using HPCToolkit. We can transform the CCT generated by HPCToolkit into a flat profile by specifying the column on which to apply the groupby operation and the function to use for aggregation. In this case, we use `sum` to get the total time spent in each function.

Load imbalance: When program developers run their code on a large number of MPI processes, load imbalance across processes is often a scaling bottleneck. Hatchet makes it extremely easy to

	nid	time	time (inc)		nid	time	time (inc)		nid	time	time (inc)
name				module				file			
<unknown file> [kripke]:0	17234	1.825282e+08	1.825282e+08	'Kripke/build-mvapich2.3/kripke	14366	1.825802e+08	5.847993e+08	own file> [kripke]	50314	3.651083e+08	1.802709e+09
Kernel_3d_DGZ::scattering	60	7.669936e+07	7.896253e+07	4/gcc-4.9.3/hpctoolkit-develop- oolkit/ext-libs/libmonitor.so.0.0.0	2512	0.000000e+00	1.918548e+08	86_64/memset.S	26693	6.148785e+06	1.229496e+07
Kernel_3d_DGZ::LTimes	30	5.010439e+07	5.240528e+07	/usr/lib64/ld-2.17.so	9676	0.000000e+00	9.340625e+02	libpsm2.so.2.1	783495	1.041419e+06	3.537456e+06
Kernel_3d_DGZ::LPlusTimes	115	4.947707e+07	5.104498e+07	/usr/lib64/libc-2.17.so	37970	0.000000e+00	7.150550e+06	/malloc/malloc.c	180270	9.252864e+05	1.859379e+06
Kernel_3d_DGZ::sweep	981	5.018862e+06	5.018862e+06	/usr/lib64/libdl-2.17.so	4427	0.000000e+00	2.804062e+02	sm/handlemem.c	10844	4.230814e+05	6.440107e+05
memset.S:99	3773	3.168982e+06	3.168982e+06	/usr/lib64/libpsm2.so.2.1	433252	0.000000e+00	2.496037e+06	src/mpidi_calls.c	17239	2.530799e+05	2.607438e+05
memset.S:101	3970	2.120895e+06	2.120895e+06	/usr/lib64/libpthread-2.17.so	2679	0.000000e+00	4.674375e+02	src/psm_queue.c	73291	2.066301e+05	1.599298e+06
Grid_Data::particleEdit	1201	1.131266e+06	1.249157e+06	c-4.9.3/lib64/libstdc++.so.6.0.20	14945	0.000000e+00	3.898480e+05	il/pt2pt/testany.c	5704	1.746973e+05	1.605984e+06
<unknown file> [libpsm2.so.2.1]:0	324763	9.733415e+05	9.733415e+05	npiler/lib/intel64_lin/libintc.so.5	1215	0.000000e+00	9.357812e+01	yscall-template.S	13787	6.691503e+04	1.338301e+05
memset.S:98	3767	6.197776e+05	6.197776e+05	intel-18.0.1/lib/libmpi.so.12.1.1	126726	0.000000e+00	7.962225e+06	[libmpi.so.12.1.1]	24482	2.587261e+04	5.043725e+04

```

1 gf = GraphFrame.from_hpctoolkit('kripke')
2
3 grouped = gf.dataframe.groupby('name').sum() # replace 'name' with 'module' or 'file'

```

Figure 9: Generating a flat profile in Hatchet using the groupby functionality of pandas. Traditional tools create a flat profile based on function names or callsite labels. In Hatchet, you can choose any categorical column to group by: name of the function (left figure), load module (middle figure), or file (right figure).

study load imbalance across processes or threads at the per-node granularity (call site or function level). A typical metric to measure imbalance is to look at the ratio of the maximum and average time spent in a code region across all processes. If the maximum-to-average ratio is high, it represents heavy imbalance. On the other hand, if the ratio is close to one, that signifies a well-balanced code.

Figure 10 shows the code for calculating an imbalance metric in an execution profile. We perform a `drop_index_levels` operation on the `GraphFrame` in two different ways: by providing mean as a function in one case and max as the function to another copy of the `DataFrame`. This generates two `DataFrames`, one containing the average time spent in each node, and the other containing the maximum time spent in each node by any process. If we divide the corresponding columns of the two `DataFrames` and look at the nodes with the highest value of the max-to-average ratio, we have located the nodes with highest imbalance. The figure shows all the nodes in LULESH with an imbalance greater than 2.0.

6.3 Comparing Execution Profiles

Another important task in parallel performance analysis is comparing the performance of an application on two different thread counts or process counts. This typically entails generating two sets of profiles on the different process counts in question and then comparing them in a GUI. Most tool GUIs do not provide automated ways to compare multiple datasets. As a result, in most cases the user manually goes over the two datasets in two instances of the tool to look for areas of the tree or graph where the performance looks different. This can be extremely cumbersome, inefficient and in many cases, ineffective. The filter, squash and subtract operations provided by the Hatchet API can be extremely powerful in comparing profiling datasets from two executions.

On-node scaling: In the first example, we ran LULESH in two modes: on a single core of a node and using 27 cores on a node. We generated profiles for the two executions and we wanted to identify

node	name	nid	time	time (inc)	imbalance
LagrangeNodal	LagrangeNodal	3.0	2.242594e+06	2.593621e+07	2.494720
main	main	0.0	1.106013e+05	5.357208e+07	2.161845
CalcForceForNodes	CalcForceForNodes	4.0	1.036399e+06	2.369361e+07	2.142526
CalcQForElems	CalcQForElems	16.0	3.351894e+06	6.649351e+06	2.037651
CalcEnergyForElems	CalcEnergyForElems	22.0	1.571996e+06	2.807323e+06	2.013174
CalcPressureForElems	CalcPressureForElems	23.0	1.235327e+06	1.235327e+06	2.005437

```

1 gf1 = GraphFrame.from_caliper('lulesh-512cores')
2 gf2 = gf1.copy()
3
4 gf1.drop_index_levels(function=np.mean)
5 gf2.drop_index_levels(function=np.max)
6
7 gf1.dataframe['imbalance']
8     = gf2.dataframe['time'].div(gf1.dataframe['time'])

```

Figure 10: Load imbalance within a single execution is derived by calculating the mean and maximum values of a metric at each node across all MPI processes or threads and then dividing the two values for each node.

the most important code regions in LULESH with respect to increase in time as one scales on node. Figure 11 presents the code to do such analysis. We read in the two profiles into two different `GraphFrames` and do a subtract (after dropping the additional index level using a mean function). The `GraphFrame` returned by the subtraction shows the nodes that have the largest increase in execution time. Although the nodes with the largest absolute time in the two cases is `CalcHourglassControlForElems`, the largest increase in time between the two executions happens on the `TimeIncrement` node (time shown in red in the right graph visualization).

Note that we do not need to visualize the graph to find the nodes with bottlenecks, especially when the graphs are very large. We can analyze the `DataFrame` of the `GraphFrame` returned by the subtract operation. Each column in the new `DataFrame` contains

Figure 11: The subtract operation in Hatchet enables comparing execution profiles. In this figure, the left graph is subtracted from the middle graph to obtain the right graph. When we sort the nodes in the right graph by time, we can easily identify the biggest offenders.

Figure 12: Hatchet makes it easy to extract the calls in a particular library, MPI for example, using the filter operation, and then to compare the extracted sub-graphs using the subtract operation. In the example above, we can easily identify which specific MPI_Send calls take more time when we scale from 27 to 512 cores.

the results of row-wise subtraction of the two input DataFrames. Sorting the columns in the new DataFrame by decreasing time can quickly identify the most problematic nodes.

Multi-node scaling: In a similar scenario, a user might be interested in comparing two executions that use a different number of MPI processes. Let's say that the user is interested in finding the difference in times spent in different MPI routines by call site. We can do this also using the Hatchet API and a few lines of code.

Figure 12 shows the code for this analysis. We read in the two datasets of LULESH, and filter them both on the name column by matching the names against ^MPI. After the filtering operation, we squash the DataFrames to generate GraphFrames that just contain the MPI calls from the original datasets. We can now subtract

the squashed datasets to identify the biggest offenders. In the figure, we observe that as we scale from 27 to 512 cores, the largest time increase is in MPI_Allreduce. As we can see, the graph and DataFrame objects in Hatchet and the powerful pandas API can help in simplifying complex performance analysis tasks, which would have possibly taken many man-hours in another tool.

Finally, we demonstrate the use of Hatchet for comparing several datasets to study the weak scaling behavior of an application. We ran LULESH from 1 to 512 cores on third powers of some numbers (a requirement of the application). We read in all the datasets into Hatchet, and for each dataset, we use a few lines of code to filter the regions where the code spends most of the time. We then use the pandas' pivot and plot operations to generate a stacked bar chart that shows how the time spent in different regions of LULESH

```

1 datasets = glob.glob('lulesh*.json')
2 datasets.sort()
3
4 dataframes = []
5 for dataset in datasets:
6     gf = GraphFrame.from_caliper(dataset)
7     gf.drop_index_levels()
8
9     num_pes = re.match('(.*)-(\d+)(.*)', dataset).group(2)
10    gf.dataframe['pes'] = num_pes
11    filtered_gf = gf.filter(lambda x: x['time'] > 1e6)
12    dataframes.append(filtered_gf.dataframe)
13
14 result = pd.concat(dataframes)
15 pivot_df = result.pivot(index='pes', columns='name', values='time')
16 pivot_df.loc[:, :].plot.bar(stacked=True, figsize=(10,7))

```

Figure 13: We read in eight LULESH caliper datasets and create a GraphFrame for each. We then filter the datasets to focus on the most time-consuming regions. For plotting, we concatenate all the DataFrames into one while storing a key that identifies the number of processes, and then use pivot to rearrange the data in a format more suitable for pandas’ plot function. The resulting stacked bar chart is shown on the right.

changes as the code scales (Figure 13). This case study demonstrates the combined potential of Hatchet and pandas in making data analytics quick and convenient for the HPC user. We believe no other performance tool provides the functionality to generate such information without significant time and effort.

7 RELATED WORK

There are many profilers that can display call path or call graph profiles. Tools like Caliper [5], OpenSpeedShop [25], TAU [21], Score-P [12], and HPCToolkit [3] can all gather fine-grained execution profiles for post-mortem performance analysis. CallFlow [18], hpcviewer [17] in HPCToolkit, TAU, and flame graphs [9] are visualization tools specifically for CCTs. All of these tools have their own format for storing the collected data, and all but Caliper and flame graphs provide a custom GUI for viewing call path profiles. Some, like TAU, can import data from other tools, but none of them offer a *programmable* interface for dealing with raw, structured profile data from parallel runs. Users must point and click to analyze the data, which can be time consuming and inflexible for large datasets or custom analyses.

Many performance tools provide facilities to store performance data in a database and to apply machine learning and other data analysis tools to it. PerfExplorer [11] provides a database, a GUI analysis environment, and the PerfDMF [10] data format. OpenSpeedShop has an internal SQL database used by the GUI to load parts of performance datasets. However, all of these tools predate the popularization of data analysis frameworks like R [19] and pandas [15, 16], and they do not provide rich APIs for manipulating data. TauDB, part of PerfDMF, provides language bindings for exploring datasets, but it does not provide the in-memory query or aggregation capabilities that modern frameworks have. All “queries” in these tools must be written in SQL, with a fixed schema, and handed off directly to the backend database. There is no in-memory DataFrame or abstraction layer as we have leveraged in Hatchet. The closest related work to Hatchet is likely *differential profiling*. Early work [14, 20] showed the utility of subtracting similar or scaled call trees to

pinpoint performance issues. This work was improved upon by techniques for scaling analysis implemented in HPCToolkit [23, 24]. HPCToolkit provides facilities for calculating derived expressions from performance metrics on call trees *within* the GUI, and this can be used to scale and subtract columns in the hpcviewer GUI. However, the usage model is cell-based like a spreadsheet; it is not fully programmable or easily integrated with other frameworks.

Likely the most scalable existing call path *visualizer* is HPCTraceViewer [22], which provides visualizations of call paths over time, MPI ranks, and threads in parallel codes. This tool and Libra [6] are the closest analogs to the per-MPI-rank analyses in this paper. Again, though, these are GUI tools and they do not provide the flexibility to easily script new analyses or to easily query, filter, aggregate, and squash profile data in an indexed DataFrame as Hatchet does. Typically, the available analyses are manually selected through drop down menus or some other user-interface, and there is limited flexibility for customization.

With Hatchet, we provide a common data model for representing structured profiles from today’s HPC tools. We provide a means to index a DataFrame by structured attributes, such as nodes in a call tree or call graph, and Hatchet builds on the widely used pandas data analysis framework, and all of the plotting and analysis libraries that can be used with it. Hatchet is not a closed-universe tool; it provides a *canonical* representation of profile data and can read data from many existing tools. If Hatchet users need to analyze data from a new measurement tool, they can do so without modifying their analysis scripts, and without learning a new format, new API, or new GUI. We advocate the use of existing *measurement* tools with Hatchet for analysis, in order to achieve more automated, reproducible results.

8 CONCLUSION

Analyzing performance and connecting performance degradation to parts of the code is important to guide application developers in their performance optimization efforts. Large parallel applications with tens to thousands of lines of codes are difficult to analyze.

Additionally, performance profiles of such applications can have hundreds of thousands of call sites or nodes in a dynamic execution profile. Most existing tools fall short in allowing users to programmatically analyze performance data.

In this paper, we presented Hatchet, a Python-based library leveraging the powerful API of data analysis tools, such as pandas to analyze structured profiling data. Since pandas does not support structured data indexed by nodes in a graph, Hatchet provides a hierarchical index to support indexing DataFrame rows by nodes in the graph. Hatchet provides a canonical data model that enables representing and analyzing different types of performance data.

Leveraging many DataFrame operations and adding its own, Hatchet simplifies many common performance analysis tasks on structured profiling data. Using case studies, we demonstrated that Hatchet provides an easy way to perform many complex tasks on parallel profiles by writing a few lines of code. These tasks include, 1) identifying regions or call sites with the most load imbalance across MPI processes or threads, 2) filtering datasets by a metric or library/function names to focus on subsets of data; and 3) easily handling and analyzing multi-rank, profile data from multiple executions. We expect that Hatchet will make HPC performance analysis quicker, easier, and more effective.

In the future, we plan to add a query language for the Hatchet user to specify expressions for filtering a graph. The user should be able to specify for example, select all the nodes whose load module is X and a descendant is an MPI routine. We believe that with a language to specify different ways to dissect and prune graphs and trees, Hatchet could become even more powerful in the terms of the kinds of analyses it can support.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-CONF-772402).

REFERENCES

- [1] [n.d.]. *Hydrodynamics Challenge Problem*, Lawrence Livermore National Laboratory. Technical Report LLNL-TR-490254.
- [2] [n.d.]. Kripke. <https://codesign.llnl.gov/kripke.php>.
- [3] Laksono Adhianto, Sinchan Banerjee, Mike Fagan, Mark Krentel, Gabriel Marin, John Mellor-Crummey, and Nathan R Tallent. 2010. HPCToolkit: Tools for Performance Analysis of Optimized Parallel Programs. *Concurrency and Computation: Practice and Experience* 22, 6 (2010), 685–701.
- [4] D. Böehme, D. Beckingsale, and M. Schulz. 2017. Flexible Data Aggregation for Performance Profiling. In *2017 IEEE International Conference on Cluster Computing (CLUSTER)*. 419–428. <https://doi.org/10.1109/CLUSTER.2017.34>
- [5] David Boehme, Todd Gamblin, David Beckingsale, Peer-Timo Bremer, Alfredo Gimenez, Matthew LeGendre, Olga Pearce, and Martin Schulz. 2016. Caliper: Performance Introspection for HPC Software Stacks. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC '16)*. IEEE Computer Society, Article 47, 11 pages. <http://dl.acm.org/citation.cfm?id=3014904.3014967> LLNL-CONF-699263.
- [6] Todd Gamblin, Bronis R. de Supinski, Martin Schulz, Robert J. Fowler, and Daniel A. Reed. 2008. Scalable Load-Balance Measurement for SPMD Codes. In *Supercomputing 2008 (SC'08)*. Austin, Texas. <http://www.cs.unc.edu/~tgamblin/pubs/wavelet-sc08.pdf> LLNL-CONF-406045.
- [7] Susan L. Graham, Peter B. Kessler, and Marshall K. McKusick. 1982. Gprof: A call graph execution profiler. *SIGPLAN Not.* 17, 6 (1982), 120–126.
- [8] Brendan Gregg. [n.d.]. Flame Graphs. <https://github.com/brendangregg/FlameGraph>.
- [9] Brendan Gregg. 2015. Flame graphs. Online. http://www.brendangregg.com/Slides/FreeBSD2014_FlameGraphs.pdf.
- [10] Kevin Huck, Allen D. Malony, R. Bell, L. Li, and A. Morris. 2005. PerfDMF: Design and implementation of a parallel performance data management framework. In *International Conference on Parallel Processing (ICPP'05)*.
- [11] Kevin A. Huck and Allen D. Malony. 2005. PerfExplorer: A Performance Data Mining Framework For Large-Scale Parallel Computing. In *Supercomputing 2005 (SC'05)*. Seattle, WA, 41.
- [12] Andreas Knüpfer, Christian Rössel, Dieter an Mey, Scott Biersdorff, Kai Diethelm, Dominic Eschweiler, Markus Geimer, Michael Gerndt, Daniel Lorenz, Allen Malony, Wolfgang E. Nagel, Yury Oleynik, Peter Philippen, Pavel Saviankou, Dirk Schmidl, Sameer Shende, Ronny Tschüter, Michael Wagner, Bert Wesarg, and Felix Wolf. 2012. Score-P: A Joint Performance Measurement Run-Time Infrastructure for Periscope, Scalasca, TAU, and Vampir. In *Tools for High Performance Computing 2011*, Holger Brunst, Matthias S. Müller, Wolfgang E. Nagel, and Michael M. Resch (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 79–91.
- [13] AJ Kunen, TS Bailey, and PN Brown. 2015. KRIPKE-A massively parallel transport mini-app. *Lawrence Livermore National Laboratory (LLNL)*, Livermore, CA, Tech. Rep (2015).
- [14] P. E. McKenney. 1995. Differential profiling. In *MASCOTS '95. Proceedings of the Third International Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*. 237–241. <https://doi.org/10.1109/MASCOT.1995.378681>
- [15] Wes McKinney. 2010. Data Structures for Statistical Computing in Python. In *Proceedings of the 9th Python in Science Conference*, Stéfan van der Walt and Jarrod Millman (Eds.). 51 – 56.
- [16] Wes McKinney. 2017. *Python for Data Analysis: Data Wrangling with Pandas, NumPy, and IPython*. O'Reilly Media. <https://www.amazon.com/Python-Data-Analysis-Wrangling-IPython/dp/1491957662?SubscriptionId=AKIAIOBINVZYXZQZ2U3A&tag=chimbiori05-20&linkCode=xm2&camp=2025&creative=165953&creativeASIN=1491957662>
- [17] J. Mellor-Crummey, R. Fowler, and G. Marin. 2002. HPCView: A tool for top-down analysis of node performance. *The Journal of Supercomputing* 23 (2002), 81–101.
- [18] H. T. Nguyen, L. Wei, A. Bhate, J. Gamblin, D. Boehme, M. Schulz, K. Ma, and P. Bremer. 2016. VIPACT: A Visualization Interface for Analyzing Calling Context Trees. In *2016 Third Workshop on Visual Performance Analysis (VPA)*. 25–28. <https://doi.org/10.1109/VPA.2016.009>
- [19] R Core Team. 2013. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org/>
- [20] Martin Schulz and Bronis R. de Supinski. 2007. Practical Differential Profiling. In *Euro-Par 2007 Parallel Processing*, Anne-Marie Kermarrec, Luc Bougé, and Thierry Priol (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 97–106.
- [21] S. Shende and A. D. Malony. 2006. The TAU Parallel Performance System. *International Journal of High Performance Computing Applications* 20, 2 (2006), 287–311. <https://doi.org/10.1177/1094342006064482>
- [22] N. Tallent, J. Mellor-Crummey, M. Franco, R. Landrum, and L. Adhianto. 2011. Scalable Fine-grained Call Path Tracing.
- [23] Nathan R. Tallent, Laksono Adhianto, and John M. Mellor-Crummey. 2010. Scalable Identification of Load Imbalance in Parallel Executions Using Call Path Profiles.
- [24] Nathan R. Tallent, John M. Mellor-Crummey, Laksono Adhianto, Michael W. Fagan, and Mark Krentel. 2011. Diagnosing performance bottlenecks in emerging petascale applications.
- [25] The Open|SpeedShop Team. [n.d.]. Open|SpeedShop for Linux. <http://www.openspeedshop.org>
- [26] David Wheeler. 2012. SLOccount. <http://www.dwheeler.com/sloccount>

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

We ran Kripke and LULESH on Quartz (an Intel Xeon cluster at LLNL). Each application was run on multiple node counts and used 32 cores per node.

ARTIFACT AVAILABILITY

Software Artifact Availability: All author-created software artifacts are maintained in a public repository under an OSI-approved license.

Hardware Artifact Availability: There are no author-created hardware artifacts.

Data Artifact Availability: All author-created data artifacts are maintained in a public repository under an OSI-approved license.

Proprietary Artifacts: None of the associated artifacts, author-created or otherwise, are proprietary.

List of URLs and/or DOIs where artifacts are available:

Hatchet: <https://github.com/LLNL/hatchet>
Kripke: <https://github.com/LLNL/Kripke>
LULESH: <https://github.com/LLNL/LULESH>
Caliper: <https://github.com/LLNL/Caliper>
HPCToolkit: <https://github.com/HPCToolkit/hpctoolkit>

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: Quartz at LLNL

Compilers and versions: Intel 18.0.1

Libraries and versions: MVAPICH 2.2

Output from scripts that gathers execution environment information.

```
SLURM_NODELIST=quartz46
LMOD_FAMILY_COMPILER_VERSION=18.0.1
SLURM_CHECKPOINT_IMAGE_DIR=/var/slurm/checkpoint
SLURM_JOB_NAME=sh
MANPATH=/usr/tce/packages/cmake/cmake-3.5.2/share/man:
↳ /usr/tce/packages/mvapich2/mvapich2-2.3-intel-
↳ 18.0.1/man:/usr/tce/packages/intel/intel-18.0.1/
↳ man/common:/usr/tce/man:/usr/share/lmod/lmod/sha
↳ re/man:/usr/tce/packages/dotkit/dotkit/man:/usr/
↳ man:/usr/share/man:/usr/local/man:/usr/X11R6/man
↳ :/usr/lib64/mvapich/default/man
CENTER_JOB_ID=quartz-2607761
SCOREP_MPI_ENABLE_GROUPS=ENV,P2P,COLL,XNONBLOCK
HOSTNAME=quartz1532
GUESTFISH_INIT=\e[1;34m
```

```
_ModuleTable003_=cHJvcFQ9e30sWyJzdGFja0RlcHRoIl09MSx_
↳ bInN0YXR1cyJdPSJhY3RpdmUiLFsidXNlck5hbWUiXT0ibXZ_
↳ hcGljaDIvMi4zIix9LHRleGxpdmU9e1siZm4iXT0iL3Vzci9_
↳ 0Y2Uvbw9kdWx1ZmlsZXMvQ29yZS90ZXhsaXZlLzIwMTYubHV_
↳ hIixbImZ1bGx0YW11Ii09InRleGxpdmUvMjAxNiIsWyJsb2F_
↳ kT3JkZXIiXT0zLHByb3BUPXt9LlFsc3RhY2tEZXB0aCJdPTE_
↳ sWyJzdGF0dXMiXT0iYWNoaXZlIixbInVzZXJ0YW11Ii09InR_
↳ leGxpdmUvMjAxNiIsfSx9LG1wYXR0QT17Ii91c3IvdGNlL21_
↳ vZHVzZWZpbGVzL01QSS9tdmFwaWNoMi8yLjMiLCIvdXNyL3R_
↳ jZS9tb2R1bGVmaWx1cy9NUeKvaW50ZWwvMTguMC4xL212YXB_
↳ pY2gylzIuMyIsIi91c3IvdGNlL21vZHVzZWZpbGVzL0Nv
SLURM_TOPOLOGY_ADDR=quartz46
SLURMD_NODENAME=quartz46
SLURM_PPIO_PROCESS=0
INTEL_LICENSE_FILE=/usr/tce/packages/intel/intel-18.
↳ 0.1/compilers_and_libraries_2018.1.163/linux/Lic
↳ enses/license.client.intel.lic
SLURM_SRUN_COMM_PORT=33016
SHELL=/bin/bash
TERM=xterm
__LMOD_REF_COUNT_MODULEPATH=/usr/tce/modulefiles/MPI
↳ /mvapich2/2.3.1;/usr/tce/modulefiles/MPI/intel/1
↳ 8.0.1/mvapich2/2.3.1;/usr/tce/modulefiles/Compil
↳ er/intel/18.0.1.1;/collab/usr/global/tools/modul
↳ efiles/toss_3_x86_64_ib/Core:1;/usr/tce/modulefi
↳ les/Core:1;/usr/apps/modulefiles:1;/usr/share/mo
↳ dulefiles/Linux:1;/usr/share/modulefiles/Core:1;
↳ /usr/share/lmod/lmod/modulefiles/Core:1
SLURM_JOB_QOS=normal
HISTSIZE=GOTCHA
WISECONFIGDIR=/usr/share/wise2/
SLURM_PTY_WIN_ROW=45
I_MPI_FABRICS=shm:tmi
SSH_CLIENT=76.102.78.232 62251 22
TMPDIR=/var/tmp/USER
MODULEPATH_ROOT=/usr/share/modulefiles
LMOD_SYSTEM_DEFAULT_MODULES=StdEnv
SLURM_TOPOLOGY_ADDR_PATTERN=node
LMOD_PKG=/usr/share/lmod/lmod
TMI_CONFIG=/etc/tmi.conf
SLURM_CPU_BIND_VERBOSE=quiet
QTDIR=/usr/lib64/qt-3.3
QTINC=/usr/lib64/qt-3.3/include
LMOD_VERSION=7.8.16
SSH_TTY=/dev/pts/37
LC_ALL=C
__LMOD_REF_COUNT_LOADEDMODULES=intel/18.0.1.1;mvapic
↳ h2/2.3.1;texlive/2016.1;StdEnv:1;cmake/3.5.2.1
SLURM_CPU_BIND_LIST=0xFFFFFFFFFFFFFFFF
HFI_NO_CPUAFFINITY=1
SLURM_NNODES=1
USER=USER
```



```

LS_COLORS=no=00:fi=00:di=01;34:ln=00;36:pi=40;33:so=
  01;35:do=01;35:bd=40;33;01:cd=40;33;01:or=41;33;
  01:ex=00;32:*.cmd=00;32:*.exe=01;32:*.com=01;32:
  *.bat=01;32:*.btm=01;32:*.dll=01;32:*.tar=00;31:
  *.tbz=00;31:*.tgz=00;31:*.rpm=00;31:*.deb=00;31:
  *.arj=00;31:*.taz=00;31:*.lzh=00;31:*.zip=00;31:
  *.zoo=00;31:*.Z=00;31:*.Z=00;31:*.gz=00;31:*.bz2
  =00;31:*.tb2=00;31:*.tz2=00;31:*.tbz2=00;31:*.av
  i=01;35:*.bmp=01;35:*.fli=01;35:*.gif=01;35:*.jp
  g=01;35:*.jpeg=01;35:*.mng=01;35:*.mov=01;35:*.m
  pg=01;35:*.pcx=01;35:*.pbm=01;35:*.pgm=01;35:*.p
  ng=01;35:*.ppm=01;35:*.tga=01;35:*.tif=01;35:*.x
  bm=01;35:*.xpm=01;35:*.dl=01;35:*.gl=01;35:*.wmv
  =01;35:*.aiff=00;32:*.au=00;32:*.mid=00;32:*.mp3
  =00;32:*.ogg=00;32:*.voc=00;32:*.wav=00;32:

LMOD_sys=Linux
LD_LIBRARY_PATH=/usr/tce/packages/mvapich2/mvapich2-
  2.3-intel-18.0.1/lib:/usr/tce/packages/intel/intel-
  el-18.0.1/lib/intel64
_dk_shell=bash
ENV=/g/g92/USER/.bashrc
PFTP_CONFIG_FILENAME=/etc/pftp_config
IPATH_NO_CPUAFFINITY=1
SLURM_STEP_NUM_NODES=1
SLURM_JOBID=2607761
HOST_GRP=linux
_ModuleTable004_bXBpbGVyL2ludGVsLzE4LjAuMSIsIi9jb2x
  sYWVvdXNyL2dsb2JhcC90b29scy9tb2R1bGVmaWxlcY90b3N
  zXzNfeDg2XzY0X2liL0NvcmluLzRjZS9tb2R1bGV
  maWxlcY9Db3JlIiwil3Vzci9hcHBzL21vZHVhZS9tb2R1bGV
  iL3Vzci9zaGFyZS9tb2R1bGVmaWxlcY9MaW51eCIsIi9lc3I
  vc2hhcmUvbW9kdWx1Zm1sZXNvQ29yZSIsIi9lc3Ivc2hhcmU
  vbG1vZC9sbW9kL21vZHVhZS9tb2R1bGVmaWxlcY9MaW51eC
  0ZW1CYXNlTVBBVEgiXT0iL3Vzci90Y2UvbW9kdWx1Zm1sZXN
  vQ29yZTovdXNyL2FwcHMvbm9kdWx1Zm1sZXN6L3Vzci9zaGF
  yZS9tb2R1bGVmaWxlcY9MaW51eDovdXNyL3NoYXJlL21vZHV
  sZWZpbGVzL0NvcmluLzRjZS9tb2R1bGVmaWxlcY9MaW51eD
SRUN_DEBUG=3
__LMOD_REF_COUNT__LMFILES_=/usr/tce/modulefiles/Core
  /intel/18.0.1.lua:1;/usr/tce/modulefiles/Compile
  r/intel/18.0.1/mvapich2/2.3.lua:1;/usr/tce/modul
  efiles/Core/texlive/2016.lua:1;/usr/tce/modulefi
  les/Core/StdEnv.lua:1;/usr/tce/modulefiles/Core/
  cmake/3.5.2.lua:1
SCOREP_REDUCE_PROBE_TEST=1
MV2_ENABLE_AFFINITY=0
SLURM_LAUNCH_NODE_IPADDR=192.168.70.8
DK_UEQRU=1
SLURM_STEP_ID=0
GUESTFISH_PS1=\[\e[1;32m\]>>>f>\[\e[0;31m\]
LMOD_PREPEND_BLOCK=normal
LMOD_FAMILY_MPI_VERSION=2.3
IBV_FORK_SAFE=1
SLURM_TASKS_PER_NODE=36

```

```

PATH=/g/g92/USER/bin/toss_3_x86_64_ib/scorep-3.0/bin
: /usr/lib/jvm/java-1.7.0-oracle-1.7.0.101.x86_64
/bin:/usr/tce/packages/cmake/cmake-3.5.2/bin:/usr
r/tce/packages/texlive/texlive-2016/2016/bin/x86
_64-linux:/usr/tce/packages/mvapich2/mvapich2-2.
3-intel-18.0.1/bin:/usr/tce/packages/intel/intel
-18.0.1/bin:/usr/tce/bin:/usr/lib64/qt-3.3/bin:/u
sr/local/bin:/usr/bin:/usr/local/sbin:/usr/sbin
MAIL=/var/spool/mail/USER
_ModuleTable001_=X01vZHV5ZVRhYmx1Xz17WyJNVHJZLnNpb24
iXT0ZLFsiY19yZWJ1aWxkVGl0ZSJdPWZhbnN1LFsiY19zaG9
ydFRpbWUixT1mYWxzZSxkZXB0aFQ9e30sZmFtaWx5PXBtImN
vbXBpbGVyIl09ImludGVsIixbIm1waSJdPSJtdmFwaWN0MiI
sfSxtVD17U3RkRW52PXBtImZuIl09Ii91c3IvdGN1L21vZHV
sZWZpbGVzL0NvcmlUvU3RkRW52Lmx1YSIsWyJmdWxsTmFtZSJ
dPSJtdGRFbnYiLFI5bG9hZE9yZGVyIl09NCxwcm9wVD17fSx
bInN0YW9nRGVwdGgiXT0wLFIsc3RhdmVzIl09ImFjdGl2ZSI
sWyJ1c2VyTmFtZSJdPSJtdGRFbnYiLH0sY21ha2U9e1siZm4
iXT0iL3Vzci90Y2UvbW9kdWx1ZmlsZXMvQ29yZS9jbWFRZS8
zLjUuMi5sdWEiLFI5ZnVsbnVsbWUixT0iY21ha2UvMy41
SLURM_STEP_LAUNCHER_PORT=33016
LCSCHEDCLUSTER=quartz
SLURM_WORKING_CLUSTER=quartz:equartz187:6817:8192
_=/usr/bin/env
SLURM_JOB_ID=2607761
IPATH_NO_BACKTRACE=1
LMOD_SETTARG_CMD=:
PWD=/g/g92/USER
INPUTRC=/etc/inputrc
_dk_r1=-1
SLURM_STEPID=0
SLURM_JOB_USER=USER
_LMFILES=/usr/tce/modulefiles/Core/intel/18.0.1.lua
: /usr/tce/modulefiles/Compiler/intel/18.0.1/mvap
ich2/2.3.lua:/usr/tce/modulefiles/Core/texlive/2
016.lua:/usr/tce/modulefiles/Core/StdEnv.lua:/usr
r/tce/modulefiles/Core/cmake/3.5.2.lua
JAVA_HOME=/usr/lib/jvm/java-1.7.0-oracle-1.7.0.101.x
86_64
SLURM_SRUN_COMM_HOST=192.168.70.8
EDITOR=/bin/vi
SLURM_CPU_BIND_TYPE=mask_cpu:
LANG=en_US.UTF-8
MODULEPATH=/usr/tce/modulefiles/MPI/mvapich2/2.3:/usr
r/tce/modulefiles/MPI/intel/18.0.1/mvapich2/2.3:
/usr/tce/modulefiles/Compiler/intel/18.0.1:/coll
ab/usr/global/tools/modulefiles/toss_3_x86_64_ib
/Core:/usr/tce/modulefiles/Core:/usr/apps/module
files:/usr/share/modulefiles/Linux:/usr/share/mo
dulefiles/Core:/usr/share/lmod/lmod/modulefiles/
Core
SLURM_UMASK=0007
SLURM_PTY_WIN_COL=80
GUESTFISH_OUTPUT=\e[0m
KDEDIRS=/usr

```

Hatchet: Pruning the Overgrowth in Parallel Profiles

```
LOADED_MODULES=intel/18.0.1:mvapich2/2.3:texlive/2016_
↳ :StdEnv:cmake/3.5.2
_ModuleTable_Sz_=5
PSM2_CLOSE_GRACE_INTERVAL=0
SLURM_JOB_UID=38725
SLURM_NODEID=0
SLURM_SUBMIT_DIR=/g/g92/USER
DK_SUBNODE=bash/toss_3_x86_64_ib:bash:toss_3_x86_64_
↳ ib:.
LMOD_CMD=/usr/share/lmod/lmod/libexec/lmod
_ModuleTable005_=bW9kdWx1Zm1sZXNMQ29yZSIscFQ==
SLURM_TASK_PID=76003
SLURM_CPUS_ON_NODE=36
KRB5CCNAME=FILE:/tmp/krb5cc_38725_yn32WW
HISTCONTROL=erasedups
ENVIRONMENT=INTERACTIVE
_dk_inuse=lcinit.0 dkcomplete.1 reuse.1
SSH_ASKPASS=
SLURM_PROCID=0
SLURM_JOB_NODELIST=quartz46
VIADEV_NUM_RDMA_BUFFER=4
HOME=/g/g92/USER
SHLVL=4
SLURM_PTY_PORT=40512
__LMOD_REF_COUNT_PATH=/usr/tce/packages/cmake/cmake-
↳ 3.5.2/bin:1;/usr/tce/packages/texlive/texlive-20
↳ 16/2016/bin/x86_64-linux:1;/usr/tce/packages/mva
↳ pich2/mvapich2-2.3-intel-18.0.1/bin:1;/usr/tce/p
↳ ackages/intel/intel-18.0.1/bin:1;/usr/tce/bin:1;
↳ /usr/lib64/qt-3.3/bin:1;/usr/local/bin:1;/usr/bi
↳ n:1;/usr/local/sbin:1;/usr/sbin:1
SLURM_LOCALID=0
GREP_OPTIONS=--color=auto
SLURM_CLUSTER_NAME=quartz
SLURM_JOB_CPUS_PER_NODE=36
_ModuleTable002_=LjIiLFsibG9hZE9yZGVyIl09NSxwcm9wVDI_
↳ 7fSxbIn0YWNrRGVwdGgiXT0wLFsic3RhdHVzIl09ImFjdG1_
↳ 2ZSIscWYj1c2VyTmFtZSJdPSJjbWFrZSIscFQ29yZSIscFQ2_
↳ mbiJdPSIvdXNyL3RjZS9t2R1bGVmaWxlcY9Db3JlL2ludGV_
↳ sLzE4LjAuMS5sdWEiLFsiZnVsbE5hbWUiXT0iaW50ZWwvMTg_
↳ uMC4xIixbImxvYWRPcmRlciJdPTEscHJvcFQ9e30sWyJzdGF_
↳ ja0RlCHRoIl09MSxbIn09YXR1cyJdPSJhY3RpdmUiLFsidXN_
↳ lck5hbWUiXT0iaW50ZWwvMT09bXZhcGljaDI9e1siZm4iXT0_
↳ iL3Vzci90Y2Uvbw9kdWx1Zm1sZXNMQ29tcGlsZXIvaW50ZWw_
↳ vMTguMC4xL212YXBpY2gyLzIuMy5sdWEiLFsiZnVsbE5hbWU_
↳ iXT0ibXZhcGljaDIvMi4zIixbImxvYWRPcmRlciJdPTIs
SLURM_JOB_GID=38725
SLURM_SUBMIT_HOST=quartz1532
SLURM_GTIDS=0
SLURM_JOB_PARTITION=pdebug
BASH_ENV=/usr/share/lmod/lmod/init/bash
LOGNAME=USER
LMOD_arch=x86_64
PSM2_CLOSE_TIMEOUT=1
CVS_RSH=ssh
QTLIB=/usr/lib64/qt-3.3/lib
```

```
VIADEV_DEFAULT_TIME_OUT=22
SLURM_STEP_NUM_TASKS=1
SLURM_JOB_ACCOUNT=lc
SSH_CONNECTION=76.102.78.232 62251 134.9.54.10 22
_dk_isatty=1
XDG_DATA_DIRS=/g/g92/USER/.local/share/flatpak/expor
↳ ts/share:/var/lib/flatpak/exports/share:/usr/loc
↳ al/share:/usr/share
SLURM_JOB_NUM_NODES=1
SYS_TYPE=toss_3_x86_64_ib
MODULESHOME=/usr/share/lmod/lmod
LESSOPEN=||/usr/bin/lesspipe.sh %s
LMOD_SETTARG_FULL_SUPPORT=no
__LMOD_REF_COUNT_LD_LIBRARY_PATH=/usr/tce/packages/m
↳ vapich2/mvapich2-2.3-intel-18.0.1/lib:1;/usr/tce
↳ /packages/intel/intel-18.0.1/lib/intel64:1
SCOREP_ENABLE_TRACING=1
SCOREP_ENABLE_PROFILING=0
SLURM_STEP_TASKS_PER_NODE=1
LMOD_FULL_SETTARG_SUPPORT=no
LMOD_FAMILY_COMPILER=intel
SLURM_STEP_NODELIST=quartz46
__LMOD_REF_COUNT_INTEL_LICENSE_FILE=/usr/tce/package
↳ s/intel/intel-18.0.1/compilers_and_libraries_201
↳ 8.1.163/linux/Licenses/license.client.intel.lic:1
QT_PLUGIN_PATH=/usr/lib64/kde4/plugins:/usr/lib/kde4
↳ /plugins
LMOD_DIR=/usr/share/lmod/lmod/libexec
__LMOD_REF_COUNT_MANPATH=/usr/tce/packages/cmake/cma
↳ ke-3.5.2/share/man:1;/usr/tce/packages/mvapich2/
↳ mvapich2-2.3-intel-18.0.1/man:1;/usr/tce/package
↳ s/intel/intel-18.0.1/man/common:1;/usr/tce/man:1
↳ ;/usr/share/lmod/lmod/share/man:1;/usr/tce/packa
↳ ges/dotkit/dotkit/man:1;/usr/man:1;/usr/share/ma
↳ n:1;/usr/local/man:1;/usr/X11R6/man:1;/usr/lib64
↳ /mvapich/default/man:1
SLURM_CPU_BIND=quiet,mask_cpu:0xFFFFFFFFFFFFFFFF
GUESTFISH_RESTORE=\e[0m
HISTFILE=/g/g92/USER/.bash_history
DK_ROOT=/usr/tce/packages/dotkit/dotkit
DK_NODE=/usr/global/tools/dotkit
LMOD_COLORIZE=yes
LMOD_FAMILY_MPI=mvapich2
SCOREP_INSTALL=/g/g92/USER/bin/toss_3_x86_64_ib/scor
↳ ep-3.0
VIADEV_DEFAULT_MIN_RNR_TIMER=25
BASH_FUNC_module()=() { eval "$($LMOD_CMD bash "$@")
↳ && eval "${LMOD_SETTARG_CMD:-} -s sh)
}
BASH_FUNC_ml()=() { eval "$($LMOD_DIR/ml_cmd "$@")
}
+ lsb_release -a
LSB Version: :core-4.1-amd64:core-4.1-noarch:
↳ cxx-4.1-amd64:cxx-4.1-noarch:desktop-4.1-amd64:d
↳ esktop-4.1-noarch:languages-4.1-amd64:languages-
↳ 4.1-noarch:printing-4.1-amd64:printing-4.1-noarch
```

```

Distributor ID: RedHatEnterpriseServer
Description: Red Hat Enterprise Linux Server
↳ release 7.6 (Maipo)
Release: 7.6
Codename: Maipo
+ uname -a
Linux quartz46 3.10.0-957.5.1.3chaos.ch6.x86_64 #1
↳ SMP Tue Feb 26 14:19:40 PST 2019 x86_64 x86_64
↳ x86_64 GNU/Linux
+ lscpu
Architecture: x86_64
CPU op-mode(s): 32-bit, 64-bit
Byte Order: Little Endian
CPU(s): 72
On-line CPU(s) list: 0-71
Thread(s) per core: 2
Core(s) per socket: 18
Socket(s): 2
NUMA node(s): 2
Vendor ID: GenuineIntel
CPU family: 6
Model: 79
Model name: Intel(R) Xeon(R) CPU E5-2695 v4
↳ @ 2.10GHz
Stepping: 1
CPU MHz: 2101.000
CPU max MHz: 2101.0000
CPU min MHz: 1200.0000
BogoMIPS: 4190.31
Virtualization: VT-x
L1d cache: 32K
L1i cache: 32K
L2 cache: 256K
L3 cache: 46080K
NUMA node0 CPU(s): 0-17,36-53
NUMA node1 CPU(s): 18-35,54-71
Flags: fpu vme de pse tsc msr pae mce
↳ cx8 apic sep mtrr pge mca cmov pat pse36 clflush
↳ dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx
↳ pdpe1gb rdtscp lm constant_tsc arch_perfmon pebs
↳ bts rep_good nopl xtopology nonstop_tsc
↳ aperfmperf eagerfpu pni pclmulqdq dtes64 monitor
↳ ds_cpl vmx smx est tm2 ssse3 sdbg fma cx16 xtpr
↳ pdcm pcid dca sse4_1 sse4_2 x2apic movbe popcnt
↳ tsc_deadline_timer aes xsave avx f16c rdrand
↳ lahf_lm abm 3dnowprefetch epb cat_l3 cdp_l3
↳ intel_ppin intel_pt ssbd ibrs ibpb stibp
↳ tpr_shadow vnmi flexpriority ept vpid fsgsbase
↳ tsc_adjust bmi1 hle avx2 smep bmi2 erms invpcid
↳ rtm cqm rdt_a rdseed adx smap xsaveopt cqm_llc
↳ cqm_occup_llc cqm_mbm_total cqm_mbm_local dtherm
↳ ida arat pln pts spec_ctrl intel_stibp flush_l1d
+ cat /proc/meminfo
MemTotal: 131583488 kB
MemFree: 128808820 kB
MemAvailable: 128655336 kB

```

```

Buffers: 4424 kB
Cached: 449428 kB
SwapCached: 0 kB
Active: 168724 kB
Inactive: 420108 kB
Active(anon): 143804 kB
Inactive(anon): 134944 kB
Active(file): 24920 kB
Inactive(file): 285164 kB
Unevictable: 11396 kB
Mlocked: 11412 kB
SwapTotal: 0 kB
SwapFree: 0 kB
Dirty: 0 kB
Writeback: 0 kB
AnonPages: 146152 kB
Mapped: 42240 kB
Shmem: 136272 kB
Slab: 583932 kB
SReclaimable: 99620 kB
SUnreclaim: 484312 kB
KernelStack: 19552 kB
PageTables: 4488 kB
NFS_Unstable: 0 kB
Bounce: 0 kB
WritebackTmp: 0 kB
CommitLimit: 65791744 kB
Committed_AS: 474932 kB
VmallocTotal: 34359738367 kB
VmallocUsed: 1212060 kB
VmallocChunk: 34291281968 kB
HardwareCorrupted: 0 kB
AnonHugePages: 10240 kB
CmaTotal: 0 kB
CmaFree: 0 kB
HugePages_Total: 0
HugePages_Free: 0
HugePages_Rsvd: 0
HugePages_Surp: 0
Hugepagesize: 2048 kB
DirectMap4k: 1457424 kB
DirectMap2M: 66603008 kB
DirectMap1G: 68157440 kB
+ inxi -F -c0
./collect_env.sh: line 14: inxi: command not found
+ lsblk -a
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda 8:0 0 26.1G 1 disk
└─rootdev 253:0 0 26.1G 1 mpath
sdb 8:16 0 26.1G 1 disk
└─rootdev 253:0 0 26.1G 1 mpath
+ lsscsi -s
[10:0:0:0] disk LIO-ORG FILEIO 4.0
↳ /dev/sda 28.0GB
[11:0:0:0] disk LIO-ORG FILEIO 4.0
↳ /dev/sdb 28.0GB

```

Hatchet: Pruning the Overgrowth in Parallel Profiles

```
+ module list
++ /usr/share/lmod/lmod/libexec/lmod bash list

Currently Loaded Modules:
  1) intel/18.0.1   3) texlive/2016   5) cmake/3.5.2
  2) mvapich2/2.3   4) StdEnv

+ eval 'MODULEPATH=/usr/tce/modulefiles/MPI/mvapich2'
+ /2.3:/usr/tce/modulefiles/MPI/intel/18.0.1/mvapi
+ ch2/2.3:/usr/tce/modulefiles/Compiler/intel/18.0
+ .1:/collab/usr/global/tools/modulefiles/toss_3_x
+ 86_64_ib/Core:/usr/tce/modulefiles/Core:/usr/app
+ s/modulefiles:/usr/share/modulefiles/Linux:/usr/
+ share/modulefiles/Core:/usr/share/lmod/lmod/modu
+ lefiles/Core;
export 'MODULEPATH;' '_ModuleTable001_=X01vZHVSVZRH
+ mx1Xz17WyJNVHZlcnNpb24iXT0zLFsiY19yZWJ1aWxkVGltZ
+ SJdPWZhbHN1LFsiY19zaG9ydFRpbWUiXT1mYWxzZXsXkZXB0a
+ FQ9e30sZmFtaWx5PXBtbImNvbXBpbGVyIi09ImIudGVsIixbI
+ m1waSJdPSJtdmFwaWNoMiIsfSxtVD17U3RkRW52PXBtbImZuI
+ l09Ii91c3IvdGN1L21vZHVSVZWZpbGVzL0NvcmlUvU3RkRW52L
+ mx1YSIsWyJmdWxsTmFtZSJdPSJtdGRFbnYiLFsibG9hZE9yZ
+ GvYiI09NCxwcm9wVD17fSxbInN0YWNrRGVwdGgiXT0wLFsic
+ 3RhHVzIi09ImFjdG12ZSIsWyJ1c2VyTmFtZSJdPSJtdGRFb
+ nYiLH0sY21ha2U9e1siZm4iXT0iL3Vzci90Y2UvbW9kdWx1Z
+ mlsZXMvQ29yZS9jbWFrZS8zLjUuMi5sdWEiLFsiZnVsbE5hb
+ WUiXT0iY21ha2UvMy41;'
export '_ModuleTable001_'; '_ModuleTable002_=LjIiLFs
+ ibG9hZE9yZGVyIi09NSxwcm9wVD17fSxbInN0YWNrRGVwdGg
+ iXT0wLFsic3RhHVzIi09ImFjdG12ZSIsWyJ1c2VyTmFtZSJd
+ dPSJjbWFrZSIsfSxpbmRlbnB17WyJmbiJdPSIvdXNyL3RjZS9
+ tb2R1bGVmaWxlcY9Db3JlL2ludGVsLzE4LjAuMS5sdWEiLFs
+ iZnVsbE5hbWUiXT0iaW50ZWwvMTguMC4xIixbImxvYWRPcmR
+ lciJdPTEscHJvcFQ9e30sWyJzdGFja0RlcHRoIi09MSxbInN
+ 0YXR1cyJdPSJhY3RpdmlUfSsidXNlck5hbWUiXT0iaW50ZWw
+ iLH0sbXZhcGljaDI9e1siZm4iXT0iL3Vzci90Y2UvbW9kdWx
+ lZmlsZXMvQ29tGlsZXIvaW50ZWwvMTguMC4xL212YXBpY2g
+ yLzIuMy5sdWEiLFsiZnVsbE5hbWUiXT0ibXZhcGljaDIvMi4
+ zIixbImxvYWRPcmRlciJdPTIs;'
export '_ModuleTable002_'; '_ModuleTable003_=cHJvcFQ
+ 9e30sWyJzdGFja0RlcHRoIi09MSxbInN0YXR1cyJdPSJhY3R
+ pdmUilFsidXNlck5hbWUiXT0ibXZhcGljaDIvMi4zIix9LHR
+ leGxpdmU9e1siZm4iXT0iL3Vzci90Y2UvbW9kdWx1ZmlsZXM
+ vQ29yZS90ZXhsaXZlLzIwMTYubHVhIixbImZ1bGx0YW11Ii0
+ 9InRleGxpdmUvMjAxNiIsWyJsb2FkT3JkZXIiXT0zLHByb3B
+ UPXt9LFsic3RhY2tEZXB0aCJdPTEsWyJzdGF0dXMiXT0iYWw
+ 0aXZlIixbInVzZXJOYV11Ii09InRleGxpdmUvMjAxNiIsfSx
+ 9LG1wYXR0QT17Ii91c3IvdGN1L21vZHVSVZWZpbGVzL01QSS9
+ tdmFwaWNoMi8yLjMiLCIvdXNyL3RjZS9tb2R1bGVmaWxlcY9
+ NUEkvaW50ZWwvMTguMC4xL212YXBpY2gyLzIuMyIsIi91c3I
+ vdGN1L21vZHVSVZWZpbGVzL0Nv;'
```

```
export '_ModuleTable003_'; '_ModuleTable004_=bXBpbGV
+ yL2ludGVsLzE4LjAuMSIsIi9jb2xsYWlvdXNyL2dsb2JhbC9
+ 0b29scy9tb2R1bGVmaWxlcY90b3NzXzNfeDg2XzY0X2liL0N
+ vcmUilCIdXNyL3RjZS9tb2R1bGVmaWxlcY9Db3JlIiwiL3V
+ zci9hcHBzL21vZHVSVZWZpbGVzIiwiL3Vzci9zaGFyZS9tb2R
+ 1bGVmaWxlcY9MaW51eCIsIi91c3Ivc2hhcmUvbW9kdWx1Zml
+ sZXMvQ29yZSIsIi91c3Ivc2hhcmUvbG1vZC9sbW9kL21vZHV
+ sZWZpbGVzL0NvcmlUilH0sWyJzeXN0ZW1CYXN1TVBBVEgiXT0
+ iL3Vzci90Y2UvbW9kdWx1ZmlsZXMvQ29yZTovdXNyL2FwcHM
+ vbW9kdWx1ZmlsZXM6L3Vzci9zaGFyZS9tb2R1bGVmaWxlcY9
+ MaW51eDovdXNyL3NoYXJlL21vZHVSVZWZpbGVzL0NvcmlU6L3V
+ zci9zaGFyZS9sbW9kL2xtb2Qv;'
export '_ModuleTable004_';
+ '_ModuleTable005_=bW9kdWx1ZmlsZXMvQ29yZSIsfQ==;'
export '_ModuleTable005_'; '_ModuleTable_Sz_=5;'
export '_ModuleTable_Sz_';
++ MODULEPATH=/usr/tce/modulefiles/MPI/mvapich2/2.3:
+ /usr/tce/modulefiles/MPI/intel/18.0.1/mvapich2/2
+ .3:/usr/tce/modulefiles/Compiler/intel/18.0.1:/c
+ ollab/usr/global/tools/modulefiles/toss_3_x86_64
+ _ib/Core:/usr/tce/modulefiles/Core:/usr/apps/mod
+ ulefiles:/usr/share/modulefiles/Linux:/usr/share
+ /modulefiles/Core:/usr/share/lmod/lmod/modulefil
+ es/Core
++ export MODULEPATH
++ _ModuleTable001_=X01vZHVSVZRHymx1Xz17WyJNVHZlcnNp
+ b24iXT0zLFsiY19yZWJ1aWxkVGltZSJdPWZhbHN1LFsiY19z
+ aG9ydFRpbWUiXT1mYWxzZXsXkZXB0aFQ9e30sZmFtaWx5PXBt
+ ImNvbXBpbGVyIi09ImIudGVsIixbIm1waSJdPSJtdmFwaWNo
+ MiIsfSxtVD17U3RkRW52PXBtbImZuIi09Ii91c3IvdGN1L21v
+ ZHVSVZWZpbGVzL0NvcmlUvU3RkRW52Lmx1YSIsWyJmdWxsTmFt
+ ZSJdPSJtdGRFbnYiLFsibG9hZE9yZGVyIi09NCxwcm9wVD17
+ fSxbInN0YWNrRGVwdGgiXT0wLFsic3RhHVzIi09ImFjdG12
+ ZSIsWyJ1c2VyTmFtZSJdPSJtdGRFbnYiLH0sY21ha2U9e1si
+ Zm4iXT0iL3Vzci90Y2UvbW9kdWx1ZmlsZXMvQ29yZS9jbWFr
+ ZS8zLjUuMi5sdWEiLFsiZnVsbE5hbWUiXT0iY21ha2UvMy41
++ export _ModuleTable001_
++ _ModuleTable002_=LjIiLFsibG9hZE9yZGVyIi09NSxwcm9w
+ VD17fSxbInN0YWNrRGVwdGgiXT0wLFsic3RhHVzIi09ImFj
+ dG12ZSIsWyJ1c2VyTmFtZSJdPSJjbWFrZSIsfSxpbmRlbnB17
+ WyJmbiJdPSIvdXNyL3RjZS9tb2R1bGVmaWxlcY9Db3JlL2lu
+ dGVsLzE4LjAuMS5sdWEiLFsiZnVsbE5hbWUiXT0iaW50ZWwv
+ MTguMC4xIixbImxvYWRPcmRlciJdPTEscHJvcFQ9e30sWyJz
+ dGFja0RlcHRoIi09MSxbInN0YXR1cyJdPSJhY3RpdmlUfSsi
+ dXNlck5hbWUiXT0iaW50ZWwvMTguMC4xL212YXBpY2gyLzIu
+ My5sdWEiLFsiZnVsbE5hbWUiXT0ibXZhcGljaDIvMi4zIixb
+ ImxvYWRPcmRlciJdPTIs
++ export _ModuleTable002_
```



```

++ _ModuleTable003_=chJvcFQ9e30sWyJzdGFja0RlcHRoIl09
↪ MSxbInN0YXR1cyJdPSJhY3RpdmUiLFsidXNlck5hbWUiXT0i
↪ bXZhcG1jaDIvMi4zIix9LHRleGxpdmU9e1siZm4iXT0iL3Vz
↪ ci90Y2UvbW9kdWx1Zm1sZXNvQ29yZS90ZXhsaXZlZlIwMTYu
↪ bHVhIixbImZ1bGx0YWI1I09InRleGxpdmUvMjAxNiIsWyJs
↪ b2Fkt3JkZXIiXT0zLHByb3BUPXt9LFsic3RhY2tEZXB0aCJd
↪ PTEsWyJzdGF0dXMiXT0iYWNoaXZlIixbInVzZXJOYW1lI09
↪ InRleGxpdmUvMjAxNiIsfSx9LG1wYXR0QT17Ii91c3IvdGNl
↪ L21vZHVhZWZpbGVzL01QSS9tdmFwaWNoMi8yLjMiLCIvdXNy
↪ L3RjZS9tb2R1bGVmaWxlcY9NUeKvaW50ZWwvMTguMC4xL212
↪ YXBpY2gyLzIuMyIsIi91c3IvdGNlL21vZHVhZWZpbGVzL0Nv
++ export _ModuleTable003_
++ _ModuleTable004_=bXBpbGVyL2ludGVsLzE4LjAuMSIsIi9j
↪ b2xsYWVldXNyL2dsb2JhbC90b29scy9tb2R1bGVmaWxlcY90
↪ b3NzXzNfeDg2XzY0X2li0NvcmUiLCIvdXNyL3RjZS9tb2R1
↪ bGVmaWxlcY9Db3JlIiwil3Vzci9hcHBzL21vZHVhZWZpbGVz
↪ Iiwil3Vzci9zaGFyZS9tb2R1bGVmaWxlcY9MaW51eCIsIi91
↪ c3Ivc2hhcmUvbW9kdWx1Zm1sZXNvQ29yZSIsIi91c3Ivc2hh
↪ cmUvbG1vZC9sbW9kL21vZHVhZWZpbGVzL0NvcmUiLH0sWyJz
↪ eXN0ZW1CYXNlTVBBVEgiXT0iL3Vzci90Y2UvbW9kdWx1Zm1s
↪ ZXNvQ29yZTovdXNyL2FwcHMvbw9kdWx1Zm1sZXNvL3Vzci9z
↪ aGFyZS9tb2R1bGVmaWxlcY9MaW51eDovdXNyL3NoYXJlL21v
↪ ZHVhZWZpbGVzL0NvcmU6L3Vzci9zaGFyZS9sbW9kL2xtb2Qv
++ export _ModuleTable004_
++ _ModuleTable005_=bW9kdWx1Zm1sZXNvQ29yZSIsfQ==
++ export _ModuleTable005_
++ _ModuleTable_Sz_=5
++ export _ModuleTable_Sz_
++ : -s sh
+ eval
+ nvidia-smi
NVIDIA-SMI has failed because it couldn't communicate
↪ with the NVIDIA driver. Make sure that the latest
↪ NVIDIA driver is installed and running.

+ lshw -short -quiet -sanitize
+ cat
WARNING: you should run this program as super-user.
H/W path      Device      Class      Description
=====
/0
↪ system      Computer
/0/0
↪ bus         Motherboard
/0/0
↪ memory      128GiB System
↪ memory
/0/1
↪ processor   Intel(R)
↪ Xeon(R) CPU E5-2695 v4 @ 2.10GHz
/0/2
↪ processor   Intel(R)
↪ Xeon(R) CPU E5-2695 v4 @ 2.10GHz
/0/100
↪ bridge      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D DMI2
/0/100/1
↪ bridge      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D PCI Express Root Port 1
/0/100/2
↪ bridge      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D PCI Express Root Port 2
/0/100/2/0
↪ hsi0        network    Omni-Path HFI
↪ Silicon 100 Series [discrete]

```

```

/0/100/3
↪ bridge      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D PCI Express Root Port 3
/0/100/4
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
↪ 0
/0/100/4.1
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
↪ 1
/0/100/4.2
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
↪ 2
/0/100/4.3
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
↪ 3
/0/100/4.4
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
↪ 4
/0/100/4.5
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
↪ 5
/0/100/4.6
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
↪ 6
/0/100/4.7
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
↪ 7
/0/100/5
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D Map/VTd_Misc/System
↪ Management
/0/100/5.1
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D IIO Hot Plug
/0/100/5.2
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D IIO RAS/Control
↪ Status/Global Errors
/0/100/5.4
↪ generic      Xeon E7 v4/Xeon
↪ E5 v4/Xeon E3 v4/Xeon D I/O APIC
/0/100/11
↪ generic      C610/X99 series
↪ chipset SPSR
/0/100/11.4
↪ storage      C610/X99
↪ series chipset sSATA Controller [AHCI mode]
/0/100/16
↪ communication C610/X99
↪ series chipset MEI Controller #1
/0/100/16.1
↪ communication C610/X99
↪ series chipset MEI Controller #2
/0/100/1a
↪ bus          C610/X99 series
↪ chipset USB Enhanced Host Controller #2
/0/100/1c
↪ bridge      C610/X99 series
↪ chipset PCI Express Root Port #1
/0/100/1c.3
↪ bridge      C610/X99
↪ series chipset PCI Express Root Port #4
/0/100/1c.3/0
↪ display      MGA G200e
↪ [Pilot] ServerEngines (SEP1)
/0/100/1c.4
↪ bridge      C610/X99
↪ series chipset PCI Express Root Port #5

```

Hatchet: Pruning the Overgrowth in Parallel Profiles

/0/100/1c.4/0	eno1	network	I350 Gigabit	/0/1a	generic	Xeon E7 v4/Xeon
↳ Network Connection				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/100/1c.4/0.1	eno2	network	I350 Gigabit	/0/1b	generic	Xeon E7 v4/Xeon
↳ Network Connection				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/100/1d		bus	C610/X99 series	/0/1c	generic	Xeon E7 v4/Xeon
↳ chipset USB Enhanced Host Controller #1				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/100/1f		bridge	C610/X99 series	/0/1d	generic	Xeon E7 v4/Xeon
↳ chipset LPC Controller				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/100/1f.2		storage	C610/X99	/0/1e	generic	Xeon E7 v4/Xeon
↳ series chipset 6-Port SATA Controller [AHCI mode]				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/100/1f.3		bus	C610/X99 series	/0/1f	generic	Xeon E7 v4/Xeon
↳ chipset SMBus Controller				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/3		generic	Xeon E7 v4/Xeon	/0/20	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 0				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/6		generic	Xeon E7 v4/Xeon	/0/21	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 0				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/7		generic	Xeon E7 v4/Xeon	/0/22	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 0				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/8		generic	Xeon E7 v4/Xeon	/0/23	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 1				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/9		generic	Xeon E7 v4/Xeon	/0/24	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 1				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/a		generic	Xeon E7 v4/Xeon	/0/25	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 1				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/b		generic	Xeon E7 v4/Xeon	/0/26	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/c		generic	Xeon E7 v4/Xeon	/0/27	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1				↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent		
/0/d		generic	Xeon E7 v4/Xeon	/0/28	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1				↳ E5 v4/Xeon E3 v4/Xeon D R2PCIe Agent		
/0/e		generic	Xeon E7 v4/Xeon	/0/29	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R3 QPI Link Debug				↳ E5 v4/Xeon E3 v4/Xeon D R2PCIe Agent		
/0/f		generic	Xeon E7 v4/Xeon	/0/2a	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ E5 v4/Xeon E3 v4/Xeon D Ubox		
/0/10		generic	Xeon E7 v4/Xeon	/0/2b	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ E5 v4/Xeon E3 v4/Xeon D Ubox		
/0/11		generic	Xeon E7 v4/Xeon	/0/2c	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ E5 v4/Xeon E3 v4/Xeon D Ubox		
/0/12		generic	Xeon E7 v4/Xeon	/0/2d	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ E5 v4/Xeon E3 v4/Xeon D Home Agent 0		
/0/13		generic	Xeon E7 v4/Xeon	/0/2e	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ E5 v4/Xeon E3 v4/Xeon D Home Agent 0		
/0/14		generic	Xeon E7 v4/Xeon	/0/2f	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ E5 v4/Xeon E3 v4/Xeon D Home Agent 1		
/0/15		generic	Xeon E7 v4/Xeon	/0/30	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ E5 v4/Xeon E3 v4/Xeon D Home Agent 1		
/0/16		generic	Xeon E7 v4/Xeon	/0/31	generic	Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -		
/0/17		generic	Xeon E7 v4/Xeon	↳ Target Address/Thermal/RAS		
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				/0/32	generic	Xeon E7 v4/Xeon
/0/18		generic	Xeon E7 v4/Xeon	↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -		
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent				↳ Target Address/Thermal/RAS		
/0/19		generic	Xeon E7 v4/Xeon			
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent						

/0/33 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
 ↳ Channel Target Address Decoder
 /0/34 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
 ↳ Channel Target Address Decoder
 /0/35 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
 ↳ Broadcast
 /0/36 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Global Broadcast
 /0/37 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
 ↳ Channel 0 Thermal Control
 /0/38 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
 ↳ Channel 1 Thermal Control
 /0/39 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
 ↳ Channel 0 Error
 /0/3a generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
 ↳ Channel 1 Error
 /0/3b generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
 ↳ Interface
 /0/3c generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
 ↳ Interface
 /0/3d generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
 ↳ Interface
 /0/3e generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
 ↳ Interface
 /0/3f generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Target Address/Thermal/RAS
 /0/40 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Target Address/Thermal/RAS
 /0/41 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Channel Target Address
 ↳ Decoder
 /0/42 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Channel Target Address
 ↳ Decoder
 /0/43 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
 ↳ Broadcast
 /0/44 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Global Broadcast
 /0/45 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 1 -
 ↳ Channel 0 Thermal Control
 /0/46 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 1 -
 ↳ Channel 1 Thermal Control

/0/47 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 1 -
 ↳ Channel 0 Error
 /0/48 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 1 -
 ↳ Channel 1 Error
 /0/49 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
 ↳ Interface
 /0/4a generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
 ↳ Interface
 /0/4b generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
 ↳ Interface
 /0/4c generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
 ↳ Interface
 /0/4d generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 /0/4e generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 /0/4f generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 /0/50 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 /0/51 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 /0/52 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 /0/53 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
 /0/4 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
 ↳ 0
 /0/4.1 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
 ↳ 1
 /0/4.2 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
 ↳ 2
 /0/4.3 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
 ↳ 3
 /0/4.4 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
 ↳ 4
 /0/4.5 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
 ↳ 5
 /0/4.6 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
 ↳ 6
 /0/4.7 generic Xeon E7 v4/Xeon
 ↳ E5 v4/Xeon E3 v4/Xeon D Crystal Beach DMA Channel
 ↳ 7

Hatchet: Pruning the Overgrowth in Parallel Profiles

```

/0/5          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Map/VTd_Misc/System
↳ Management
/0/5.1        generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D IIO Hot Plug
/0/5.2        generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D IIO RAS/Control
↳ Status/Global Errors
/0/5.4        generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D I/O APIC
/0/54         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 0
/0/55         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 0
/0/56         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 0
/0/57         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 1
/0/58         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 1
/0/59         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D QPI Link 1
/0/5a         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
/0/5b         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
/0/5c         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R3 QPI Link 0/1
/0/5d         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R3 QPI Link Debug
/0/5e         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/5f         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/60         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/61         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/62         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/63         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/64         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/65         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/66         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/67         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/68         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/69         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent

```

```

/0/6a         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/6b         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/6c         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/6d         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/6e         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/6f         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/70         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/71         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/72         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/73         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/74         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/75         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/76         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Caching Agent
/0/77         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R2PCIe Agent
/0/78         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D R2PCIe Agent
/0/79         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Ubox
/0/7a         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Ubox
/0/7b         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Ubox
/0/7c         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Home Agent 0
/0/7d         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Home Agent 0
/0/7e         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Home Agent 1
/0/7f         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Home Agent 1
/0/80         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
↳ Target Address/Thermal/RAS
/0/81         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
↳ Target Address/Thermal/RAS
/0/82         generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
↳ Channel Target Address Decoder

```



```

/0/83          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
↳ Channel Target Address Decoder
/0/84          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Broadcast
/0/85          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Global Broadcast
/0/86          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
↳ Channel 0 Thermal Control
/0/87          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
↳ Channel 1 Thermal Control
/0/88          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
↳ Channel 0 Error
/0/89          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 0 -
↳ Channel 1 Error
/0/8a          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface
/0/8b          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface
/0/8c          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface
/0/8d          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 0/1
↳ Interface
/0/8e          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Target Address/Thermal/RAS
/0/8f          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Target Address/Thermal/RAS
/0/90          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Channel Target Address
↳ Decoder
/0/91          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Channel Target Address
↳ Decoder
/0/92          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Broadcast
/0/93          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Global Broadcast
/0/94          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 1 -
↳ Channel 0 Thermal Control
/0/95          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 1 -
↳ Channel 1 Thermal Control
/0/96          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 1 -
↳ Channel 0 Error

```

```

/0/97          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Memory Controller 1 -
↳ Channel 1 Error
/0/98          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface
/0/99          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface
/0/9a          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface
/0/9b          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D DDRIO Channel 2/3
↳ Interface
/0/9c          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
/0/9d          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
/0/9e          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
/0/9f          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
/0/a0          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
/0/a1          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
/0/a2          generic      Xeon E7 v4/Xeon
↳ E5 v4/Xeon E3 v4/Xeon D Power Control Unit
/0/a3          system      PnP device
↳ PNP0b00
/0/a4          system      PnP device
↳ PNP0c02
/0/a5          communication PnP device
↳ PNP0501
WARNING: output may be incomplete or inaccurate, you
↳ should run this program as super-user.

```