Towards A Programmable Analysis and Visualization Framework for Interactive Performance Analytics

Tanzima Z. Islam Department of Computer Science Texas State University tanzima@txstate.edu Alexis Ayala, Quentin Jensen Department of Computer Science Western Washington University {ayalaa2,jensenq}@wwu.edu Khaled Z. Ibrahim Lawrence Berkeley National Laboratory kzibrahim@lbl.gov

Abstract—Understanding the performance characteristics of applications in modern HPC environments is becoming more challenging due to the increase in the architectural and programming complexities. HPC software developers rely on sources such as hardware counters and event traces to infer performance problems while focusing on designing mitigation strategies. A large number of in-house tools exist in the community, which indicate replicated effort. This paper presents a customizable framework for analyzing performance measurements and visualizing through a web-based interactive dashboard for interactively exploring a large volume of hierarchical information. In this paper, we analyze three applications deemed important in the Exascale Computing Project (ECP) as use cases and identify as well as optimize problematic resource utilization behaviors exposed by our visualizations. This framework is a step towards a unified platform for visual identification of performance scaling bottlenecks to ease the collaboration between application developers, performance analysts, and hardware vendors.

Index Terms—Performance analysis, Machine learning, Visualization

I. INTRODUCTION

Modern systems include architectures with hundreds of components (e.g., L1, L2, memory), and their complex interactions with applications cause performance loss. These machines can log thousands of on-core and off-core events that can explain how an application interacts with the underlying system. Analyzing these events can expose potential optimization opportunities for an application. For example, a large number of page faults may cause a performance issue. Visual inspection of the collected data can help expose apparent patterns that can be useful for designing targeted analysis and optimization strategies.

While this swell of information is undoubtedly useful, this may leave performance analysts with overwhelming information to investigate. A typical performance analysis workflow includes: (1) identifying regions of interest to investigate, (2) identifying how an application utilizes the underlying hardware resources such as cache hierarchy and compute units, (3) pinpointing actual events recorded by the hardware to gain insight into performance problems, and (4) comparing performance across workloads, applications, and configurations (e.g., varying numbers of workers). Frequently, application developers team up with performance analysts to develop in-house tools written in languages such as R, Python to wrangle the vast amount of data collected to gain insights into performance problems quickly. The focus typically stays on devising optimization strategies afterward to mitigate the problems. Since the goal of the in-house analysis tools is often quick data exploration to guide optimization, these tools usually are not designed to be extendable by others. This phenomenon leads to duplicated efforts of developing similar if not the same analysis and visualization pipelines across the HPC community.

To fill this gap, we have developed an HPC performance analysis and visualization framework, DASHING, based on the principles of programmability, extendability, and interpretability of results in mind. The modular design of this Python framework enables plugging in any analysis and visualization techniques of choice. Specifically, in this paper, we present a web-based visualization dashboard for interactive exploration of hierarchical performance data. These interactive visualizations provide both coarse- and fine-grained information about the causes of a target performance metric (e.g., efficiency loss), and a quantitative ranking of these reasons so that a user can prioritize their mitigation strategies. We envision that some of these causes can be easily fixed by the application developers (i.e. domain scientists), while others may need further engagement between application developers, performance analysts with system expertise, and vendors.

Although the types of analysis and visualizations that can be supported by DASHING are practically unlimited, we demonstrate the effectiveness of this framework by addressing the following typical performance analysis questions: (1) explain the impact of resource utilization behaviors on performance, (2) identify individual hardware events that can explain scalability loss, and (3) compare two performance profiles quantitatively across applications. Here, the term resource stands for an arbitrary user-defined group of hardware performance counters. As case studies, we analyze three ECP applications – Nyx3D, IAMR, and INCFLO. Nyx3D is a compressible cosmological hydrodynamics application. IAMR and INCFLO solve incompressible Navier-Stokes equations, which are used by combustion simulations. The INCFLO application is a much smaller proxy application, developed to mimic the resource utilization behaviors of IAMR. For analysis, we extended our previous machine learning techniques [1] to provide counter importance along with resource significance in explaining the parallel efficiency loss as applications scale on a node.

While most works in the literature have focused on developing

novel analysis techniques, the community has recognized the value of visualizations to aid in exploring the plethora of collected information. Several works have developed specific visualizations for specific problem domains [2]–[6]. In contrast, we aim to develop a common platform for all researchers to easily incorporate their methodologies to serve the performance analysis needs of the community.

The contributions of this paper are:

- A novel hierarchical visualization of how the hardware utilization behaviors of applications impact its performance
- An easily extendable and modular framework, DASHING, that implements several visualization techniques both novel and widely used by the performance analysis community
- Optimization opportunities exposed by the DASHING framework for three applications important for the Exascale Computing Project (ECP)

The most notable contribution of this work is that our visualizations identified TLB misses being important contributors to the efficiency loss of Nyx3D and IAMR. We mitigated those issues by loading huge pages of 4MB on Cori. Some performance problems observed due to the use of atomic instructions may need more involved code refactoring, algorithm change or discussion with processor vendor. With the DASHING framework, we envision creating such engagement opportunities for co-design among application developers, performance analysts, and hardware vendors.

We organize the rest of the paper as follows. Section II presents a high-level description of the analysis methodologies used in this paper and Section III presents the overall workflow of the DASHING framework. Sections IV and V present the experimental setup and the observations from analyzing three ECP applications. Finally, Section VII summarizes the paper and indicates future research planned in this direction.

II. BACKGROUND

In this paper, we leverage two methodologies for calculating how significantly hardware components (e.g., cache, memory) impact the performance of applications and how well two performance profiles match. Such performance profiles can be collected by varying applications, workloads, or regions within an application. We extend the analysis techniques from our previous work [1], [7]. This section presents these two methodologies from a high-level.

Identifying important hardware events can be formulated as a feature selection problem in machine learning, where a number of features (e.g., cache miss, resource stalls) can explain the performance problems (e.g., increase in execution time) of applications. The following sections explain the target and the features used for our analysis.

Features: Hardware Counters and Resources On-node performance is typically constrained by hardware bottlenecks on the processors or in the memory system. Hence, we use hardware performance counters to explain performance loss. Hardware counters are registers in modern architectures that keep track of events such as cache misses, which impact the overall performance of applications. However, there are hundreds of such events that can be collected via PAPI [8]. To avoid overwhelming users with fine-grained information such as the count of each hardware event, we categorize the events to high-level groups based on the hardware components they affect. We refer to these high-level groups as "resources".

Target: Performance loss This paper focuses on explaining the performance loss of applications as they strongly scale on a node based on their hardware component utilization behaviors. Strong scaling is defined as increasing the number of workers (e.g., threads, processes) while keeping the total problem size fixed. In this paper, we define performance loss as the loss of efficiency as an application scales (Equation 2), since this metric impacts the final time to solution. Efficiency loss typically increases as an application uses more cores on a multi-core machine. Since every event (e.g, cache miss) measured by a hardware performance counter has a non-zero penalty, our goal is to identify those events that follow the same scaling pattern as the efficiency loss.

efficiency_loss(p) = 1 -
$$\frac{T(1)}{(T(p) \times p)}$$
 (1)

where p = number of workers (threads or processes) and T(p)= execution time when using p workers.

Resource Significance Measure (RSM) This paper leverages the machine learning approaches published in our previous work [1], [7]. We compute the impact of resource utilization on efficiency loss by building a linear predictive model with the collected hardware performance counters. In particular, we leverage the extended Orthogonal Matching Pursuit (eOMP) algorithm presented in [1]. The eOMP algorithm builds on the traditional OMP algorithm presented in [9], which selects the smallest and diverse subset of counters to describe the performance loss.

However, hardware performance counters from different hardware resources can be inter-correlated, which will not be selected by a traditional OMP algorithm. Instead, the eOMP algorithm selects τ most correlated counters (instead of one) and creates a discrete distribution based on their correlations, i.e., the higher the correlation, the higher probability to be chosen in that step. It then randomly picks a counter using the discrete distribution, computes the residual, and repeats this process until a desired sparsity (user defined) is met or the model achieves sufficient fidelity. We repeat this randomized algorithm 5000 times independently and obtain the final coefficients for counters by averaging the solutions in the ensemble. The algorithm then computes a belief metric for each of the counters, if they have a non-zero coefficient in the sparse representation. The belief value is inversely proportional to the reconstruction error of each counter. The final step combines counter-wise beliefs to their corresponding resources by obtaining the compound evidence [10].

Similarity Analysis We leverage our previous work [1], [7] to compare the resource utilization behaviors of a pair of applications. The comparative analysis approach computes a



Fig. 1: (a) Workflow of the DASHING framework. The dotted boxes represent points of interaction between users and the framework for leveraging DASHING for other use cases (b) Example configuration file with hierarchical structure.

distance between the two feature spaces defined by the set of performance counters. From high-level, the methodology computes (1) the low-dimensional projections (PCA) of the performance subspaces (*d* principle directions), (2) measures the principle angles between each of the principle directions (θ_i along direction *i*), and (3) the distance between the distributions (λ_i along direction *i*) of the collected performance data. Then the compatibility score of a resource (**r**) is computed as:

compatibility_score(r) =
$$1 - \frac{1}{d} \sum_{i} \theta_i \lambda_i$$
 (2)

III. APPROACH & FRAMEWORK

This paper focuses on presenting visualizations for analyzing application performance in a modular framework— DASHING—that is easy to extend. The visualizations are interactive and can guide users to (1) correlate the utilization of the hardware components to the performance loss of an application as it scales, (2) identify hardware events that can explain efficiency loss, and (3) enable quantitative comparison of the performances between two applications. In the context of performance analysis, we refer both application developers (i.e., domain scientists) and performance analysts with system expertise as users. Figure 1 presents the overall workflow of the DASHING framework, which mainly performs three tasks: (1) Processing (2) Analysis and (3) Visualization. This section describes the components that implement these tasks and the end-to-end use of the framework in details.

A. Processing

Figure 1a shows that a user initiates the use of the DASHING framework by writing a configuration file (step 1) where the configuration file describes (step 2) which analysis, visualization, data loader, and resource mappers to use when processing data.

Data Loader: The processing task starts with reading data from the performance files (step 3) collected from various measurement tools and converting those into a Python dic-

tionary. Performance measurement tools write data in different formats such as HDF5, Comma Separated Values (CSV), and plaintext. Figure 1c presents an example HDF5 file format that the DASHING framework expects. Each region in the HDF5 file corresponds to a group and each counter for a specific region corresponds to a dataset in that group. For each counter, each entry is present as a tuple <process_id, thread_id, repetition_id >, where process_id denotes a process id, *thread id* denotes a thread id within the process_id and repretition_id denotes a specific repetition of the region. The DASHING framework can be extended to read output from other tools as long as the data is formatted into CSV. The CSV must be formatted such that each column denotes a region and each row denotes a performance counter. Each entry pertaining to a region and a counter is a list of floats delimited by commas. These values denote the average counter value over several repetitions (if applicable) for each $< process_id, thread_id >$ used to run the application. For a purely multi-processed application, thread id can be assumed to be 1, whereas a purely multi-threaded application, the process id can be assumed to be 1. To read files in userdefined format, a performance analyst needs to write a Python function to read data to a Python dictionary and return that dictionary.

Resource Mapper: The resource_mapper module reads the event-to-group mapping (step 4) from a user provided CSV file or calls a user-defined function to map hardware events (or features) to resources (groups). In this work, we leverage the event-to-resource mapping function developed from our previous work [1] to categorize the hardware events on Cori. Table I presents the resource groups used in this work for the Intel Haswell architecture.

B. Analysis

The analysis module essentially is a wrapper that iterates through a list of analysis functions received from the processing module (step 5) written in Python. This function pointer based method enables users to perform any analysis of their choice, as long as the return value is a Python dictionary of structure $\langle key1, key2, key3 \rangle \rightarrow value$. In this paper, the tuple represents $\langle region, resource, event \rangle \rightarrow$ *importance*. The two analysis methodologies used in this work address the following research questions: (1) which regions in an application constitute the largest run time, (2) which hardware resource utilization can explain the efficiency loss curve as an application scales on a node, (3) which resource utilization behaviors match between two applications. For addressing the first question, the DASHING framework uses execution times to rank regions of interest. For addressing the second and the third questions, the DASHING framework builds on the machine learning techniques presented in our previous work [1], [7]. A high-level overview of the methods is presented in Section II.

C. Visualization

The visualization module is responsible for visually representing the output from different analysis (step 6) into a web-based dashboard. The goal of these visualizations is to present hierarchical information in an interactive manner. The DASHING framework provides detailed meaning of each performance event along with any known optimization techniques as a tool tip text for each event. This description of each event, provided through the configuration file, contains < *eventname*, *description* > pairs for each event. Since the tool will be made available on Github to the community, we envision both system analysts and architects can per take in enhancing this tool by annotating the event description files, written in the CSV format.

While the processing, analysis, and visualization tasks are performed from a command-line tool, the visualization module opens up the interactive visualizations in a web browser. Currently, users can interact with the visualizations alone (step 7), *after* an analysis is run on the entire dataset. In future, we plan to support interactivity between the web-based dashboard and the analysis module so that users can run new analysis on a selected part of the dataset.

With the goal of easy interpretability, we have created several visualizations corresponding to several commonly performed analysis tasks. This list of visualizations present information from a coarse to fine-grained level:

1) Heat map (Figures 3a, 3b)

compares the RSM scores of resources per region, giving a broad overview about which resources are heavily used by which regions. This summary visualization provides a quick way for comparing the resource utilization behaviors across all annotated regions while abstracting away a lot of details.

2) Sunburst (Figure 2[a-d])

presents a large amount of hierarchical information in a visually intuitive manner. Each section of each circle represents its relative cost to its parent. Each section is clickable, expanding that section to fill the whole chart (Figures 2b, 2c). A single sunburst chart represents an application. Moving outwards, the next circle represents the contribution of each region to the overall execution time of the application. Then, the importance of each resource that contributes to explaining the efficiency loss of its parent region. Finally, the importance of each hardware counter that contributes to its respective resource. In this fashion, a user can quickly identify what regions of code are high-cost, and then zoom in on why. This coarseto-fine grained visualization is powerful in its ability to help a user discover specific ways to optimize their code. The tool shows only significant events attributed to efficiency loss, down-selecting from hundreds of system events.

3) Interactive line chart (Figure 2e)

shows the values of hardware counters normalized within a region, as an application scales. Each line is colored by its resource, shown alongside the region's efficiency loss curve (bold black). Counters that follow the efficiency loss occurred tightly are likely to explain why the efficiency loss occurred. This visualization enables users to investigate the trends in events compared to the trend in the target variable (e.g., efficiency loss).

4) Scatter plot (Figure 3e):

showcases the comparison between two performance profiles. The X axis shows the RSM scores for the first application, and the Y axis shows how well the second application mimics the resource utilization behaviors of the first one. A lower value of compatibility indicates greater dissimilarity between the two performance profiles. This visualization was proposed in our previous work [1] and has been integrated in the DASHING framework.

Although, the Sunburst chart can present a comparative view of the resource utilization across regions, it is more suitable for a hierarchical view of the performance characteristics of a single region. In contrast, the heat map is more suitable for a summarized view of the resource utilization behaviors across regions. Hence, the DASHING framework includes both.

D. Configuration

DASHING provides a simple configuration file-based approach to execute and extend DASHING. The configuration file is written in YAML. The configuration file is organized as a hierarchy of tasks. We assume that performance data is collected for code regions (e.g., function) and is organized under a name (region name). This organization enables several regions in the same application to be collected and analyzed at once. It is written in Python 3.7, hence programmable. The command line invocation of DASHING is: python driver.py config.yaml.

Figure 1b shows an example configuration file for generating Figures 2[a-e]. The fields in red (data and tasks) describe keywords, and the items under the tasks are functions to apply on the data. To extend DASHING, users can enlist their Python function names under tasks (e.g., func1, func2). If the task pertains to a specific dataset, then it should be listed under a micro task (e.g., compute_rsm on Nyx3D data). The main task provides access to all datasets (i.e., micro tasks). Hence, tasks to analyze across several datasets should be enlisted under main. For example, the compatibility_score is compute between two regions

Resource Group on Intel Haswell	ID
Prefetcher	PRE
Floating point unit	FP
Branch instruction and prediction units	BR
L1 cache	L1
L2 cache	L2
L3 cache	L3
Offcore events	OFF
Micro operations	UOPS
Front-end	FE
Non-floating point arithmetic unit	ARITH
Atomic events	ATOMIC
Translation look aside buffer	TLB
Memory	MEM
Events to prune	UNDEFINED

TABLE I: Resource groups on Cori

in two datasets. These regions are accessed by specifying their micro task names, a "::" operator, and the region names after that (e.g., nyx::strang_hydro where Nyx3D is the micro task name and strang_hydro is the region name). In the future, we will make the configuration editing capability available via the web-based dashboard itself to reduce the adoption barrier to application developers and increase their productivity.

IV. EXPERIMENTAL SETUP

In this section, we describe the system and applications we use to demonstrate the use cases of the DASHING framework. For all applications, we used the sparsity parameter to be 3 (as was found to be sparse enough in our previous work [1]). *A. System Description*

We ran our experiments on Cori [11], a Cray XC40 supercomputer with 29.5 Petaflop/s peak performance. Cori consists of 2, 388 Intel Xeon E5-2698 (Haswell) and 9, 688 Intel Xeon Phi 7200 (Knights Landing) processors. Each Haswell compute node has two sockets, each socket has 16 cores with 36.8 Gflops/core and 128 GB memory. Each core has its own L1 and L2 caches, with 64 KB (32 KB instruction cache, 32 KB data) and 256 KB, respectively; there is also a 40-MB shared L3 cache per socket. In this paper, we present experimental results on the Intel Haswell nodes only. Table I presents the hardware resources on Cori. We collected more than 200 hardware counters and categorized them based on the open source PMU tools project [12] for Intel CPUs. We leave performance analysis on the Intel Xeon Phi processors for future work.

B. Applications

Nyx3D is an adaptive mesh, compressible cosmological hydrodynamics simulation code that leverages the AMReX framework. The application is written using a combination of Fortran 2003 and C++11 with parallelism supported by a hybrid MPI+OpenMP model [13]. Nyx3D solves equations of compressible hydrodynamics on an adaptive grid hierarchy coupled with an N-body treatment of dark matter.

IAMR is a parallel, adaptive mesh refinement (AMR) code that solves the variable-density incompressible Navier-Stokes equations. The IAMR source code can be found at [14].

INCFLO [15] is a proxy application similar in purpose to

IAMR, but has a simplified code base. INCFLO has been developed to only capture the computation and the performance characteristics of the advection, diffusion, and projection operations for solving the incompressible Navier-Stokes equations (implemented in the advance method) in its much larger counterpart, IAMR.

V. RESULTS

In this section, we present the effectiveness of DASHING visualizations using three ECP applications. We collected the hardware performance counters for the three applications Nyx3D, IAMR, and INCFLO using the Perf-dump [16] tool. Due to page limit, we only present a subset of these visualizations to demonstrate the usability of the developed framework. Specifically, we present the following visualizations for answering common questions asked during a typical performance analysis phase:

- Heat map: Comparison of the resource utilization behaviors across regions
- Sunburst: Comparison of the execution times, resource importance, and contributing hardware events as well as their importance
- Interactive line chart: Comparison of the features (e.g., hardware performance counters) with the target (e.g., execution time, efficiency loss) across various configurations (e.g., varying number of workloads, workers)
- Scatter plot: Comparison of resource utilization between two performance profiles. These profiles can be collected by varying workloads of the same application, or different regions, or even different applications.

The following section presents specific examples on how these visualizations can help both application developers and vendors identify performance optimization opportunities in applications.

A. Nyx3D

In this section, we present the resource utilization behaviors of 19 annotated regions in Nyx3D. We identified these regions through hot spot analysis using the Craypat-lite [17] tool. The main region represents the rest of the methods in the Nyx3D application that we did not explicitly annotate (i.e., not the main method itself).

Figure 2a shows the breakdown of the overall performance of Nyx3D, 2b filters by a region, and 2c filters by a resource. Tool tip text (not shown in figure) shows the numeric importance values and execution times of the regions. All regions, resources, and events are sorted based on their values (e.g., execution times, importance) and organized in counterclockwise manner.

The observations that can be made from Figure 2 are:

- strang_hydro explains 65% of the overall execution time, hence the most important region to optimize. The specific value (65%) is presented as a tool tip text when users hover over that region (not shown in figure).
- TLB is the most important resource that can explain 43% of the efficiency loss curve as Nyx3D scales on a node.
- Figure 2c shows that DTLB (Data TLB) misses cause a large number of walks that can be optimized by using huge



(e) Line charts for Nyx3D strang_hydro

Fig. 2: [a-d] Interactive graph for Nyx3D shows that ATOMIC, TLB, and Instruction stalls are the three major groups to optimize for the most expensive region, strang_hydro. [e] Interactive line chart shows the comparison of the normalized values of the features (e.g., hardware performance counters) to a target variable (e.g., the efficiency loss). Clicking on an event in the legend toggles a group.

pages.

- Figure 2d shows that using 4M huge pages on Cori indeed reduces the impact of TLB misses on the performance of Nyx3D (reduced from 11% to 3%).
- Figures 2a and 2d also show that ATOMIC instructions cause the efficiency loss to increase, which may require application developers to engage with the vendors to discuss potential optimizations.

Figure 2e shows a comparison of the normalized values of the hardware performance counters and the efficiency loss of each region as the application scales. The bold black line indicates the efficiency loss of the region as it scales on more workers, and the individual lines correspond to each important event. Hardware events are filtered by their importance to declutter the visualizations. Users can set the threshold to an arbitrary low value (e.g., zero) to control the density of information.

B. INCFLO

We ran INCFLO with multiple input sizes. The number of cells parameter was varied from $32^3(small)$, $64^3(medium)$, and $128^3(large)$. Due to the page limit, we only present a subset of the visualizations. These results clearly expose potential opportunities for optimizing the performance of the INCFLO application across multiple regions.

Figures 3a and 3b show the resource utilization behaviors across three regions annotated within the INCFLO application for the small and large problem sizes, respectively. From Figure 3a and 3b, we can observe that while TLB is an important resource across all the regions for the small input size, ATOMIC is important for the large one. From Figure 3c, we can observe that a large number of TLB store misses contribute to the increase in efficiency loss, which can be optimized by using hugepages.

C. IAMR

The DASHING framework similarly analyzes and visualizes the performance characteristics of the IAMR application. Due to page limit, we only present the comparative analysis results between IAMR (*n* cell = 32^3) and INCFLO (*n* cell = 32^3 , 64^3 , and 128^3). The INCFLO application has been developed to capture the advection, diffusion, and projection tasks for solving the incompressible Navier-Stokes equations. In this section, we only present a comparison between the performance characteristics of the advance methods (implementing all three of the tasks mentioned above) in both of these applications. Figure 3e shows the comparison between IAMR with workload 32^3 and INCFLO with workloads 32^3 and 128^3 . The X axis shows the resource significance measure for IAMR. Higher values mean a greater importance to cover. The Y axis shows the compatibility score; higher values mean a more similar behavior.

From Figure 3e we can observe that MEM is the most important resource for IAMR, however INCFLO with the small input size (Figure 3e-left) does not cover the behavior well. In this work, we deem a compatibility score of 0.8 or higher as a good match (shown using red dotted horizontal line). On the other hand, the MEM utilization behavior for INCFLO with the large input size (Figure 3e-right) covers



(d) Overall performance characteristics for (e) Comparison between IAMR and INCFLO $n_{cell} = 32^3$ and 128^3 IAMR

Fig. 3: (a) and (b) Comparison of resource utilization behaviors across regions. (c) Overall performance shows TLB is the most important resource. Application developers may use the hugepages to improve the performance of INCFLO when using a large workload. (d) MEM is the most important resource for IAMR. (e) Comparison between IAMR and INCFLO with two workloads (small and large).

that of IAMR well. This surprising observation indicates that further investigation is needed to ensure that INCFLO mimics IAMR.

VI. RELATED WORK

ParaProf [18], Tau [19], YAViT [20], HPCToolKit [21] are widely used HPC tools that provide integrated analysis and supporting visualizations for identifying performance bottleneck in applications. While these are massive tools with a large user base, they often lack the flexibility of a simple Python framework for interactive exploration of the collected data. While the Caliper [22] tool shares our aim of combining multiple data sources, the interactive visualizations presented in this work are complementary to Caliper's focus on providing a scalable data model for performance analysis tasks. Visualizations of supercomputer networks [2], network traffic [23], scalability trend [3], job placements [24], power consumption [25], flops visualization for threads [5], communication [26] are great examples of targeted visualizations that have been developed by the community. All of these works focus on developing intuitive visualizations for presenting data for a specific problem. In contrast, we envision that this work will lead to the development of a unified programmable framework for all HPC performance analysis and visualization needs.

PaScal [3] is a visualization tool for studying scalability of applications, a goal similar to our own. Performance monitoring and visualization dashboard in [27] focuses on a system for collecting cluster health statistics and visualizing a large volume of data. This system comes equipped with a large number of components that enable real time monitoring, data collection in addition to visualization. Furthermore, there are works on visualizing calling contexts [28], memory performance behaviors [4], and frameworks for projecting performance data over simulation geometry [6]. While the visualizations are useful, none of the work aims at providing a unified framework for all analysis and visualization needs of the HPC performance community. In that regard, these analysis and visualization techniques are complementary to our work.

VII. CONCLUSION

This paper presents a novel visualization framework for presenting hierarchical information about application performance and identifying critical resources limiting scaling efficiency. Our framework introduces an interactive web-based dashboard that implements several performance visualization methods that can be useful for application developers, performance analysts, and system vendors alike.

Although DASHING is written in Python, the visualization dashboard can be integrated in a Cloud instance for providing interactive exploration of the performance dataset. Such a tool can enable collaboration across multiple research labs, academics, and vendors to share data without overwhelming users with raw data. The DASHING framework supports programmability through a configuration file, extendability through function pointers, and interpretability through interactive visualizations, for which we plan to provide on-the-fly customization in the future.

We will also continue to add more visualizations to the DASHING framework, so that one can fully configure the axes from the configuration file to invoke any of the functions. We would also like to support visualization for correlations across the performance counters in a meaningful way so that the ones that are always correlated can be identified. This list can be used to filter the number of hardware performance counters that need to be collected.

ACKNOWLEDGMENT

This research used resources in Lawrence Berkeley National Laboratory and the National Energy Research Scientific Computing Center, which are supported by the U.S. Department of Energy Office of Science's Advanced Scientific Computing Research program under contract number DE-AC02-05CH11231.

REFERENCES

- T. Z. Islam, J. J. Thiagarajan, A. Bhatele, M. Schulz, and T. Gamblin, "A machine learning framework for performance coverage analysis of proxy applications," in <u>SC'16</u>: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. IEEE, 2016, pp. 538–549.
- [2] S. Cheng, W. Zhong, K. E. Isaacs, and K. Mueller, "Visualizing the topology and data traffic of multi-dimensional torus interconnect networks," IEEE Access, vol. 6, pp. 57191–57204, 2018.
- [3] F. d. A. Alex and S. Xavier-de Souza, "Pascal viewer: A tool for the visualization of parallel scalability trends," in <u>Programming and</u> Performance Visualization Tools: International Workshops, ESPT 2017 and VPA 2017, Denver, CO, USA, November 12 and 17, 2017, and ESPT 2018 and VPA 2018, Dallas, TX, USA, November 16 and 11, 2018, Revised Selected Papers. Springer, p. 250.
- [4] A. Giménez, T. Gamblin, I. Jusufi, A. Bhatele, M. Schulz, P.-T. Bremer, and B. Hamann, "Memaxes: visualization and analytics for characterizing complex memory performance behaviors," <u>IEEE transactions on</u> visualization and computer graphics, vol. 24, no. 7, pp. 2180–2193, 2017.
- [5] R. Brendel, B. Wesarg, R. Tschüter, M. Weber, T. Ilsche, and S. Oeste, "Generic library interception for improved performance measurement and insight," in <u>Programming and Performance Visualization Tools</u>. Springer, 2017, pp. 21–37.
- [6] C. Wood, M. Larsen, A. Gimenez, K. Huck, C. Harrison, T. Gamblin, and A. Malony, "Projecting performance data over simulation geometry using sosflow and alpine," in <u>Programming and Performance</u> <u>Visualization Tools</u>. Springer, 2017, pp. 201–218.
- [7] J. J. Thiagarajan, R. Anirudh, B. Kailkhura, N. Jain, T. Z. Islam, A. Bhatele, J.-s. Yeom, and T. Gamblin, "Paddle: Performance analysis

using a data-driven learning environment," in <u>2018 IEEE International</u> Parallel and Distributed Processing Symposium (IPDPS), May 2018, pp. 784–793.

- [8] P. J. Mucci, S. Browne, C. Deane, and G. Ho, "Papi: A portable interface to hardware performance counters," in <u>Proceedings of the department of</u> defense HPCMP users group conference, vol. 710, 1999.
- [9] A. M. Bruckstein, M. Elad, and M. Zibulevsky, "On the uniqueness of nonnegative sparse solutions to underdetermined systems of equations," <u>IEEE Transactions on Information Theory</u>, vol. 54, no. 11, pp. 4813– 4820, 2008.
- [10] I. Bloch, "Some aspects of dempster-shafer evidence theory for classification of multi-modality medical images taking partial volume effect into account," <u>Pattern Recognition Letters</u>, vol. 17, no. 8, pp. 905–919, 1996.
- [11] L. B. N. Laboratory, "Nersc's cray xc40 supercomputer," https://www.nersc.gov/users/computational-systems/cori/.
- [12] "Pmu-tools: A collection of tools for profile collection and performance analysis on intel cpus on top of linux perf," https://github.com/andikleen/pmu-tools.
- [13] L. B. N. Laboratory, <u>Nyx: An Adaptive Mesh, Massively-parallel</u>, Cosmological Simulation Code, https://github.com/AMReX-Astro/Nyx.
- [14] —, "Iamr: A parallel, adaptive mesh refinement (amr) code that solves the variable-density incompressible navier-stokes equations," https://github.com/AMReX-Codes/IAMR.
- [15] —, "Incflo: An amrcore-based code that solves the incompressible navier-stokes equations," https://github.com/AMReX-Codes/incflo.
- [16] L. L. N. Laboratory, Perf-dump: A tool for collecting PAPI counter values per-process, per-thread, and per-timestep, https://github.com/scalability-llnl/perf-dump.
- [17] C. Inc., <u>Craypat-lite: Cray Performance Measurement and Analysis</u> <u>Toolset</u>, <u>2016</u>, <u>https://pubs.cray.com/content/S-2376/7.0.0/cray-performance-measurement-and-analysis-tools-user-guide/craypat-lite</u>.
- [18] R. Bell, A. D. Malony, and S. Shende, "Paraprof: A portable, extensible, and scalable tool for parallel performance profile analysis," in <u>European</u> <u>Conference on Parallel Processing</u>. Springer, 2003, pp. 17–26.
- [19] S. S. Shende and A. D. Malony, "The tau parallel performance system," <u>The International Journal of High Performance Computing Applications</u>, vol. 20, no. 2, pp. 287–311, 2006.
- [20] O. Aaziz, U. Panthi, and J. Cook, "Yavit (yet another viz tool): Raising the level of abstraction in end-user hpc interactions," in <u>2017 IEEE</u> <u>International Conference on Cluster Computing (CLUSTER)</u>. IEEE, <u>2017</u>, pp. 814–817.
- [21] L. Adhianto, S. Banerjee, M. Fagan, M. Krentel, G. Marin, J. Mellor-Crummey, and N. R. Tallent, "Hpctoolkit: Tools for performance analysis of optimized parallel programs," <u>Concurrency and Computation</u>: <u>Practice and Experience</u>, vol. 22, no. 6, pp. 685–701, 2010.
- [22] D. Boehme, T. Gamblin, D. Beckingsale, P.-T. Bremer, A. Gimenez, M. LeGendre, O. Pearce, and M. Schulz, "Caliper: performance introspection for hpc software stacks," in <u>Proceedings of the International</u> <u>Conference for High Performance Computing, Networking, Storage and Analysis</u>. IEEE Press, 2016, p. 47.
- [23] H. Bhatia, N. Jain, A. Bhatele, Y. Livnat, J. Domke, V. Pascucci, and P.-T. Bremer, "Interactive investigation of traffic congestion on fat-tree networks using treescope," in <u>Computer Graphics Forum</u>, vol. 37, no. 3. Wiley Online Library, 2018, pp. 561–572.
- [24] A. Bhatele, K. Mohror, S. H. Langer, and K. E. Isaacs, "There goes the neighborhood: performance degradation due to nearby jobs," in Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis. ACM, 2013, p. 41.
- [25] S. Labasan, M. Larsen, H. Childs, and B. Rountree, "Paviz: A power-adaptive framework for optimizing visualization performance." in EGPGV, 2017, pp. 1–10.
- [26] P. C. Roth, "Improved accuracy for automated communication pattern characterization using communication graphs and aggressive search space pruning," in <u>Programming and Performance Visualization Tools</u>. Springer, 2017, pp. 38–55.
- [27] T. Dang, "Visualizing multidimensional health status of data centers," in <u>Programming and Performance Visualization Tools</u>. Springer, 2017, pp. 273–283.
- [28] A. Bergel, A. Bhatele, D. Boehme, P. Gralka, K. Griffin, M.-A. Hermanns, D. Okanović, O. Pearce, and T. Vierjahn, "Visual analytics challenges in analyzing calling context trees," in <u>Programming and</u> Performance Visualization Tools. Springer, 2017, pp. 233–249.