

ytopt: Autotuning Scientific Applications for Energy Efficiency at Large Scales

Xingfu Wu, Prasanna Balaprakash, Michael Kruse, Jaehoon Koo, Brice Videau, Paul Hovland, Valerie Taylor
Argonne National Laboratory, Lemont, IL 60439

Email: {xingfu.wu,pbalapra,michael.kruse,jkoo,bvideau,hovland,vtaylor}@anl.gov

Brad Geltz, Siddhartha Jana
Intel Corporation, Hillsboro, OR 97124
Email: {brad.geltz,siddhartha.jana}@intel.com

Mary Hall
University of Utah, Salt Lake City, UT 84103
Email: mhall@cs.utah.edu

Abstract—As we enter the exascale computing era, efficiently utilizing power and optimizing the performance of scientific applications under power and energy constraints has become critical and challenging. We propose a low-overhead autotuning framework to autotune performance and energy for various hybrid MPI/OpenMP scientific applications at large scales and to explore the tradeoffs between application runtime and power/energy for energy efficient application execution, then use this framework to autotune four ECP proxy applications—XSBench, AMG, SWFFT, and SW4lite. Our approach uses Bayesian optimization with a Random Forest surrogate model to effectively search parameter spaces with up to 6 million different configurations on two large-scale production systems, Theta at Argonne National Laboratory and Summit at Oak Ridge National Laboratory. The experimental results show that our autotuning framework at large scales has low overhead and achieves good scalability. Using the proposed autotuning framework to identify the best configurations, we achieve up to 91.59% performance improvement, up to 21.2% energy savings, and up to 37.84% EDP improvement on up to 4,096 nodes.

I. INTRODUCTION

As we enter the exascale computing era, high performance, power, and energy management are key design points and constraints for any next generation of large-scale high-performance computing (HPC) systems [1]–[3]. Efficiently utilizing procured power and optimizing the performance of scientific applications under power and energy constraints are challenging for several reasons, including dynamic phase behavior, manufacturing variation, and increasing system-level heterogeneity. As the complexity of such HPC ecosystems (hardware stack, software stack, applications) continues to rise, achieving optimal performance and energy becomes a challenge. The number of tunable parameters that HPC users can configure at the system and application levels has increased significantly, resulting in a dramatically increased parameter space. Exhaustively evaluating all parameter combinations becomes very time-consuming. Therefore, autotuning for automatic exploration of the parameter space is desirable.

Autotuning is an approach that explores a search space of tunable parameter configurations of an application efficiently executed on an HPC system. Typically, one selects and evaluates a subset of the configurations on the target system and/or

uses analytical models to identify the best implementation or configuration for performance or energy within a given computational budget. However, such methods are becoming too difficult in practice because of the hardware, software, and the application complexity. Recently, the use of advanced search methods that adopt mathematical optimization methods to explore the search space in an intelligent way has received significant attention in the autotuning community. Such a strategy, however, requires search methods to efficiently navigate the large parameter search space of possible configurations in order to avoid a large number of expensive application runs to determine high-performance configurations or implementations. In this paper we propose a low-overhead machine learning (ML)-based autotuning framework to autotune four hybrid MPI/OpenMP Exascale Computing Project (ECP) proxy applications [4]—XSBench [5], SWFFT [6], AMG [7], and SW4lite [8]—to improve their performance, energy, and energy delay product (EDP) on two large-scale HPC systems: Theta [9] at Argonne National Laboratory (ANL), and Summit [10] at Oak Ridge National Laboratory (ORNL).

Traditional autotuning methods are built on heuristics that derive from automatically tuned BLAS libraries [11], experience [12]–[14], and model-based methods [15]–[18]. At the compiler level [19], ML-based methods are used for automatic tuning of the iterative compilation process [20] and tuning of compiler-generated code [21], [22]. Autotuning OpenMP codes has gone beyond loop schedules to look at parallel tasks and function inlining [23]–[26]. Recent work on leveraging Bayesian optimization to explore the parameter space search shows the potential for autotuning on CPU systems [27]–[30] and on GPU systems [31], [32]. Some recent work has used machine learning and sophisticated statistical learning methods to reduce the overhead of autotuning [33]–[36]. Most of these autotuning frameworks, however, are for autotuning on only a single or a few compute nodes using only performance as a metric.

This paper makes the following contributions.

- We propose a low-overhead autotuning framework ytopt to autotune various hybrid MPI/OpenMP applications at large scales.

- We use this ytopt framework to explore the tradeoffs between application runtime and power/energy for energy efficient application execution.
- We use this framework to autotune four ECP proxy applications, namely XSBench, AMG, SWFFT, and SW4lite, using Bayesian optimization with a Random Forest surrogate model to effectively search parameter spaces with up to 6 million different configurations.
- We demonstrate the effectiveness of our autotuning framework to tune the performance, energy, and EDP of these hybrid MPI/OpenMP applications on up to 4,096 nodes.
- The experimental results show that our proposed autotuning framework at large scales has low overhead and good scalability, providing the best configuration for the best performance, energy saving, or EDP. Using the proposed autotuning framework to identify the best configurations, we achieve up to 91.59% performance improvement, up to 21.2% energy savings, and up to 37.84% EDP improvement on up to 4,096 nodes.

The remainder of this paper is organized as follows. Section 2 discusses the background, challenges, and motivation of this study. Section 3 describes the systems and four ECP proxy applications used in this paper. Section 4 proposes our autotuning frameworks for improving performance and energy at large scales. Section 5 discusses autotuning mixed pragmas on a single node. Section 6 presents autotuning performance at large scales. Section 7 illustrates autotuning energy and EDP at large scales. Section 8 summarizes this paper.

II. BACKGROUND, CHALLENGES, AND MOTIVATION

Autotuning involves two critical requirements: (1) expression of a search space of implementations or configurations and (2) efficient navigation of the search space for identifying the optimal configuration. To address these two requirements, researchers have developed a number of autotuning frameworks that interface with application codes, libraries, and compilers to generate code variants and measure their performance [12], [16], [27], [28], [37]–[48]. They presented the expression of a collection of parameters to be tuned and their corresponding possible values, and they generated possible configurations that may or may not be valid for evaluation.

Two kinds of expressions of search space exist: vector space and tree space. Most autotuning frameworks present the search space in a vector space, that is, a fixed number of parameter knobs; these frameworks include OpenTuner [38], CLTune [42], HalideTuner [40], Orio [37], KernelTuner [44], ATF [41], [49], ytopt [27], [28], [46], GPTune [29], and Bliss [30]. The successor of HalideTuner [50] uses tree search to avoid the limitation of a vector search space but uses beam search to explore the space. ProTuner [45] further improves Halide schedule autotuning by replacing beam search with Monte Carlo tree search. The loop autotuner in Telamon also uses Monte Carlo tree search [51]. In the tradition of Halide, every level needs an assigned strategy, and a schedule where not all loops have an assigned strategy is considered

incomplete. The viability of autotuning the search space for loop transformations was demonstrated; the approach involves the straightforward representation as either a tree or a directed acyclic graph using mctree [47], [48], [52], and every loop is considered sequential until a pragma is added.

We classify autotuning frameworks into four categories: (1) enumerate all possible parameter configurations, reject invalid ones, and evaluate the valid ones [47]; (2) enumerate only valid configurations [41], [49]; (3) sample from the set of possible configurations, and reject invalid ones [23], [37], [39] during the search; and (4) sample only valid configurations, and search over them [27], [28]. The ytopt autotuning framework belongs to Category 4, which overcomes the ineffectiveness of Category 3 by generating valid samples and addresses the limitations of Categories 1 and 2, where enumerating all possible configurations can be computationally expensive for large number of parameters. However, the ytopt framework autotuned the applications only on a single computer node. Can we extend the ytopt framework to autotune MPI/OpenMP scientific applications at large scales so that we can identify the best configuration for running these applications on large-scale HPC systems efficiently? This is the main motivation of our work in this paper.

Kruse and Finkel [53] implemented a prototype of user-directed loop transformations using LLVM Clang [54] and Polly [55] with additional loop transformation pragmas such as loop reversal, loop interchange, tiling, and array packing in the ECP SOLLVE project [56]. Multiple pragmas can be composed even to the same loop and every transformation addresses different and often contradicting concerns, such as maximizing parallelism, spatial and temporal memory locality, but minimizing bandwidth and overhead. Hence there is a need to determine how to efficiently combine them to optimize an application.

In our recent work [27], [28] an autotuning framework ytopt was developed to leverage Bayesian optimization with four supervised machine learning methods—Random Forests, Gaussian Process Regression, Extra Trees, or Gradient-boosted Regression Trees—to explore the search space and identify more-promising regions, and we found the Random Forests performed the best. This autotuning framework was used to identify the optimal combination of the Clang loop pragma parameters, with the aim of improving the performance of six PolyBench benchmarks [57] and tuning the hyperparameters of a deep learning application MNIST on a single compute node.

Most of autotuning frameworks mentioned above were for autotuning on a single or a few compute nodes. Recently, new autotuning frameworks are emerging for multi-node autotuning. For example, GPTune [29] autotuned some MPI applications on up to 64 nodes with 2,048 cores with multitask learning using MPI, and Bayesian optimization was applied to increase the energy efficiency of a GPU cluster system [31]. Current large-scale HPC systems such as Theta [9] at ANL and Summit [10] at ORNL have complex system architectures and software stacks with many tunable parameters that may affect

the system performance and energy. How can we identify the best combination of these parameters for the best system performance or the lowest system energy consumption? Application developers and users often rely on these systems with the default configurations setup by the vendors to run their applications. How efficiently are these applications executed? Can we develop a low-overhead framework to autotune large-scale applications for performance and energy on large-scale HPC production systems such as Summit and Theta?

The answer to these questions is "Yes, we can." Specifically, in this paper we demonstrate a new state of practice by applying autotuning approach to optimize performance and energy of hybrid MPI/OpenMP scientific applications on up to 4,096 nodes on these systems.

III. SYSTEMS AND ECP PROXY APPLICATIONS

In this section we discuss the HPC system platforms and four ECP proxy applications [4] used in our experiments.

We conduct our experiments on the Cray XC40 Theta [9] of approximately 12 petaflops peak performance at Argonne National Laboratory and the IBM Power9 heterogeneous system Summit [10] of approximately 200 petaflops peak performance at Oak Ridge National Laboratory. In this section, we briefly describe their specifications shown in Table I.

Theta: Theta has 4,392 Cray XC40 nodes. Each node has 64 compute cores (one Intel Xeon Phi Knights Landing (KNL) 7230 with the thermal design power (TDP) of 215 W), shared L2 cache of 32 MB (1 MB L2 cache shared by two cores), 16 GB of high-bandwidth in-package memory Multi-Channel DRAM (MCDRAM), 192 GB of DDR4 RAM, and a 128 GB SSD. MCDRAM can be configured as a shared last level cache L3 (cache mode) or as a distinct NUMA node memory (flat mode) in or somewhere in between. The default memory mode is the cache mode. The Cray XC40 system uses the Cray Aries dragonfly network with user access to a Lustre parallel file system with 10 PB of capacity and 210 GB/s bandwidth.

In this work, we use GEOPM [58] to measure node energy consumption on Theta. The power sampling rate used is approximately 2 samples per second (default). We conduct all autotuning experiments in performance and energy with the cache mode. The compilers on Theta are CrayPE 2.6.5 (default) and clang 14 installed [56]. The aprun command is used to specify to ALPS (Application Level Placement Scheduler) the resources and placement parameters needed for the application at application launch on Theta.

Summit: Summit has 4,608 IBM Power System AC922 nodes. Each node contains two IBM POWER9 processors with 42 cores and six NVIDIA Volta V100 accelerators. Each node has 512 GB of DDR4 memory for use by the POWER9 processors and 96 GB of high-bandwidth memory (HBM2) for use by the accelerators. Additionally, each node has 1.6 TB of nonvolatile memory that can be used as a burst buffer. Summit is connected to an IBM Spectrum Scale filesystem providing 250 PB of storage capacity with a peak write speed of 2.5 TB/s. For each Summit node, the TDP of each Volta GPU is 300 W, and the TDP of each Power9 is 190 W. The power

TABLE I: System Platform Specifications and Tools

System Name	Cray XC40 Theta	IBM Power9 Summit
Location	Argonne National Lab	Oak Ridge National Lab
Architecture	Intel KNL	IBM Power9 + Nvidia GPU
Number of nodes	4,392	4,408
CPU cores per node	64	42
Sockets per node	1	2 for Power9; 2 for GPU sockets
CPU type and speed	Xeon Phi KNL 7230 1.30GHz	IBM Power9 4GHz
GPUs per node	None	6 Nvidia Volta GPUs
L1 cache per core	D:32KB, I:32KB	D:32KB, I:32KB
L2 cache per socket	32MB (two cores shared 1MB)	21MB (two cores shared 512KB)
L3 cache per socket	None	120MB (shared)
Threads per core	4	4
Memory per node	16GB MCDRAM, 192GB DDR4	96GB HBM2, 512GB DDR4
Network	Cray Aries Dragonfly	dual-rail EDR InfiniBand
Power tools	GEOPM, CapMC, RAPL	Nvidia-smi, NVML
TDP per socket	215W	190W/Power9; 300W/GPU
File System	Lustre PFS (210GB/s)	IBM GPFS (2.5TB/s)

consumption of each Summit node is 2,200 W. Although we use the NVIDIA System Management Interface (nvidia-smi) [59] to measure power consumption for each GPU, the power measurement for IBM Power9 is not available to the public. Therefore, we autotune only performance of HPC applications on Summit. The compilers on Summit are gcc 9.1.0 (default) and nvhpc 21.3. The jsrun command is used for managing an allocation that is provided by an external resource manager within IBM Job Step Manager (JSM) software package on Summit.

A. ECP Proxy Applications

In this section we discuss four hybrid MPI/OpenMP ECP proxy applications for our experiments: XSBench [5], SWFFT [6], AMG [7], and SW4lite [8].

1) *Weak-Scaling Applications:* We discuss the three weak-scaling ECP proxy applications XSBench, SWFFT, and AMG.

XSBench [5] is a mini-app representing a key computational kernel of the Monte Carlo neutron transport algorithm and represents the continuous energy macroscopic neutron cross section lookup kernel. It serves as a lightweight stand-in for full neutron transport applications like OpenMC [60]. This code provides a much simpler and more transparent platform for determining performance benefits resulting from a given hardware feature or software optimization. XSBench provides an MPI mode which runs the same code on all MPI ranks simultaneously with no decomposition across ranks of any kind, and all ranks accomplish the same work. It is an embarrassingly parallel implementation. It supports history-based transport (default): parallelism is expressed over independent particle histories, with each particle being simulated in a serial fashion from birth to death; and event-based transport: parallelism is instead expressed over different collision (or "event") types. XSBench is the hybrid MPI/OpenMP code written in C and supports OpenMP offload. The OpenMP offload implementation only supports the event-based transport. The problem size is large as default.

SWFFT [6] is to test the Hardware Accelerated Cosmology Code (HACC) 3D distributed memory discrete fast Fourier transform (FFT) with one forward FFT and one backward FFT. It assumes that global grid will originally be distributed

between MPI ranks using a 3D Cartesian communicator. That data needs to be re-distributed to three 2D pencil distributions in turn in order to compute the double-precision FFTs along each dimension. SWFFT is the hybrid MPI/OpenMP code written in C++ and C and requires the cubic number of MPI ranks and FFTW3 (double precision, OpenMP version) installed. We configure it as weak scaling case. The problem size is 4096x4096x4096 for 4096 MPI ranks. We also set the number of run tests 2.

AMG [7] a parallel algebraic multigrid solver for linear systems arising from problems on unstructured grids and builds linear systems for various 3-dimensional problems. Parallelism is achieved by data decomposition. AMG achieves this decomposition by simply subdividing the grid into logical X x Y x Z (in 3D) chunks of equal size. It is the hybrid MPI/OpenMP code written in C. The problem size is the 3D Laplace problem "-laplace -n 100 100 100 -P X Y Z". This will generate a problem with 1,000,000 grid points per MPI process with a domain of the size 100*X x 100*Y x 100*Z.

2) *Strong-Scaling Application:* SW4lite [8] is a bare bone version of SW4 [61], [62] (Seismic Waves, 4th order accuracy) intended for testing performance in a few important numerical kernels of SW4. SW4 implements substantial capabilities for 3-D seismic modeling with a free surface condition on the top boundary, absorbing super-grid conditions on the far-field boundaries, and an arbitrary number of point force and/or point moment tensor source terms. It uses a fourth order in space and time finite-difference discretization of the elastic wave equations in displacement formulation. The large problem LOH.1-h50 is from the SCEC (Southern California Earthquake Center) test suite [63]. It sets up a grid with a spacing h (=50) over a domain (X x Y x Z) 30000 x 30000 x 17000. It will run from time t=0 to t=9. The material properties are given by the block commands. They describe a layer on top of a half-space in the z-direction. A single moment point source is used with the time dependency being the Gaussian function. SW4lite is the hybrid MPI/OpenMP code written in C++ and Fortran90. In [64], performance and energy of SW4lite were optimized for the improved version. We use the improved version to define the parameter space for autotuning.

3) *Compiling Time for Each Application:* Table II shows the average compiling time (in seconds) for each ECP proxy application on Theta and Summit. We measured the compiling time for each application five times to get the average compiling time. We observe that the compiling time for SW4lite is 162.066 s on Theta and 58 s on Summit. This really impacts the autotuning wall-clock time in Step 4 shown in Figure 1. Because of loading the NVidia nvhpc module to compile the XSBench OpenMP offload version for using GPUs on Summit, it takes 4.645 s, which is much larger than that on Theta.

TABLE II: Compiling time (s) on Theta and Summit

System	XSBench	SWFFT	AMG	SW4lite
Theta	2.021	3.494	2.825	162.066
Summit	4.645	3.781	2.757	58.000

IV. PROPOSED AUTOTUNING FRAMEWORKS IN PERFORMANCE OR ENERGY AT LARGE SCALES

In this section we extend the ytopt autotuning framework to autotune the hybrid MPI/OpenMP applications at large scales on the ANL Theta and ORNL Summit using the metrics such as performance, energy, and EDP, where the application runtime is the primary performance metric; energy consumption captures the tradeoff between the application runtime and power consumption; and EDP captures the tradeoff between the application runtime and energy consumption.

A. Framework for Autotuning Performance at Large Scales

Figure 1 presents the framework for autotuning various hybrid MPI/OpenMP applications in performance. The application runtime is the primary metric. We analyze an application code to identify the important tunable application and system parameters (OpenMP environment variables) to define the parameter space using ConfigSpace [65] package. We use the tunable parameters to parameterize an application code as a code mold. ytopt starts with the user-defined parameter space, the code mold, and user-defined interface that specifies how to evaluate the code mold with a particular parameter configuration.

The search method within ytopt uses Bayesian optimization, where a dynamically updated Random Forest surrogate model that learns the relationship between the configurations and the performance metric, is used to balance exploration and exploitation of the search space. In the exploration phase, the search evaluates parameter configurations that improve the quality of the surrogate model, and in the exploitation phase, the search evaluates parameter configurations that are closer to the previously found high-performing parameter configurations. The balance is achieved through the use of the lower confidence bound (LCB) acquisition function that uses the surrogate models' predicted values of the unevaluated parameter configurations and the corresponding uncertainty values. The LCB acquisition function is defined in Equation 1. For the unevaluated parameter configuration x_M^i , the trained model M is used to predict a point estimate (mean value) $\mu(x_M^i)$ and standard deviation $\sigma(x_M^i)$.

$$a_{LCB}(x_M^i) = \mu(x_M^i) - \kappa\sigma(x_M^i) \quad (1)$$

where $\kappa \geq 0$ is a user-defined parameter that controls the tradeoff between exploration and exploitation. When $\kappa = 0$ for pure exploitation, a configuration with the lowest mean value is selected. When κ is set to a large value (> 1.96) for pure exploration, a configuration with large predictive variance is selected. The default value κ of is 1.96. Then the model M is updated with this selected configuration.

The iterative phase of the proposed autotuning framework in performance has the following steps:

- Step1 Bayesian optimization selects a parameter configuration for evaluation.
- Step2 The code mold is configured with the selected configuration to generate a new code.

- Step3 Based on the value of the number of threads in the configuration, the number of nodes reserved and the number of MPI ranks, aprun/jsrun command line for launching the application on the compute nodes is generated.
- Step4 The new code is compiled with other codes needed to generate an executable.
- Step5 The generated aprun/jsrun command line is executed to evaluate the application with the selected parameter configuration; the resulting application runtime is sent back to the search and recorded in the performance database.

Steps 1–5 are repeated until the maximum number of code evaluations or the wall-clock time is exhausted for the autotuning run.

In the rest of this paper, the term **ytopt processing time** includes the time spent in the parameter space search, building the surrogate model, processing the selected configuration to generate a new code and the aprun/jsrun command line, compiling the new code, launching the application, and storing the configuration and performance in the performance database (except the application runtime). We use the term **ytopt overhead** to stand for the ytopt processing time minus the application compiling time.




Fig. 1: Framework for Autotuning Performance

ytopt supports various application-tunable parameters, which impact the application performance but keep the program correctness. The application-tunable parameters can be defined as variables, pragmas, pragma clauses, a statement, or a function (piece of code). The combinations of these parameters with their ranges of values form a parameter space. This requires some knowledge about the applications and underlying systems. For instance, we define “#pragma omp parallel for” as a parameter before a loop to check how the application performance is affected with and without it. Table III presents the parameter space for the four ECP proxy applications used in our experiments, where **system param.** stands for system parameters; **application param.** stands for unique application parameters because some of them are used repeatedly in the application code; and **space size** stands for the number of configurations for the parameter space. The system parameters in this paper mainly focus on OpenMP runtime environment variables [66]: OMP_NUM_THREADS, OMP_PLACES, OMP_PROC_BIND, OMP_SCHEDULE, and the additional OMP_TARGET_OFFLOAD.

The selected application parameters which may impact performance for each application are described as follows. The two unique application parameters for XSBench are block size and additional “#pragma omp parallel for.” The five unique application parameters for XSBench-mixed (mixed Clang loop pragmas and OpenMP pragmas) are block size, Clang loop unrolling full, “#pragma omp parallel for,” and two tile sizes for 2D loop tiling. The four unique application parameters for XSBench-offload are simd, device clause, schedule for the OpenMP target pragmas, and “#pragma omp parallel for.” The one unique application parameter for SWFFT is “MPI_Barrier(CartComm);”. The three unique application parameters for AMG are “#pragma unroll(3),” “#pragma unroll(6),” and “#pragma omp parallel for.” The four unique application parameters for SW4lite are “#pragma unroll (6),” “#pragma omp parallel for,” “#pragma omp for nowait,” and “MPI_Barrier(MPI_COMM_WORLD);”. Overall, we use the parameter spaces with up to 6,272,640 configurations for our experiments.

TABLE III: Parameter Space for Each Application

ECP Proxy Apps	System param.	Application param.	Space size
XSBench	4 env. variables	2	51,840
XSBench-mixed	4 env. variables	5	6,272,640
XSBench-offload	5 env. variables	4	181,440
SWFFT	4 env. variables	1	1,080
AMG	4 env. variables	3	552,960
SW4lite	4 env. variables	4	2,211,840

B. Framework for Autotuning Energy and EDP at Large Scales

Efficiently utilizing the procured power and optimizing the performance of scientific applications under power and energy constraints are important challenges in HPC. The HPC PowerStack [1], [2], [67]—a global consortium of laboratories, vendors, and universities—has highlighted a design shift toward standardization of the HPC power-management software stack. This enables seamless integration of software solutions for managing the energy/power consumption of large-scale HPC systems. Based on the state of the art of the components available in the community for power and energy management [68]–[75], a hierarchical strawman PowerStack design [2] was proposed to manage power and energy at three levels of granularity: system level, job level, and node level. This implies the need to put in place the following incrementally: (1) define policies that govern site-level requirements, a power-aware system Resource Manager (RM) / job scheduler, a power-aware job-level manager, and a power-aware node manager; (2) define the interfaces between these layers to translate objectives at each layer into actionable items at the adjacent lower layer; and (3) drive end-to-end optimizations across different layers of the PowerStack.

In order to address these requirements, our recent work [3] (a) surveyed the high-level objectives of the existing layer-specific tuning approaches at the different layers: system (i.e., cluster), job / application, and node, (b) defined the tunable

parameters at each layer, and (c) proposed and discussed how to autotune the combination of different parameters at the distinct layers (parameter space) for an optimal solution (the smallest runtime or the lowest energy) under a system power cap.

Based on our previous work on autotuning the performance, power, and energy of applications and systems [3], [27], [28], [76], we propose a high-level end-to-end PowerStack autotuning framework for HPC systems shown in Figure 2. This diagram shows the interactions among four layers: system level, job level, node level, and application level. To the best of our knowledge, however, a practical end-to-end autotuning component is still lacking that targets all four layers for the optimal solution.

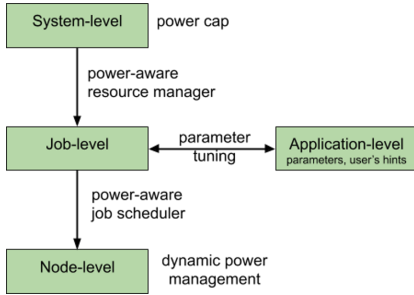


Fig. 2: PowerStack End-to-End Autotuning Framework

Our aim is to develop a practical framework to autotune all four layers of the PowerStack so that we can have a better understanding of the tunable parameters at each layer and interaction interfaces between layers and can identify potential new requirements in order to achieve energy efficiency. The process of autotuning in the layers (a) typically targets energy as the primary metric, (b) complies with the operating power constraint imposed on the layer, and (c) attempts to improve the management and orchestration of the available control parameters that affect the application and/or hardware performance. For the proposed framework, we integrate the existing job constraint-aware power/energy optimizer GEOPM (Global Extensible Open Power Manager) [58], [74] at the job and node levels and the ytopt autotuning framework at the application level.

GEOPM: An Open-source Community-driven Power Management Runtime

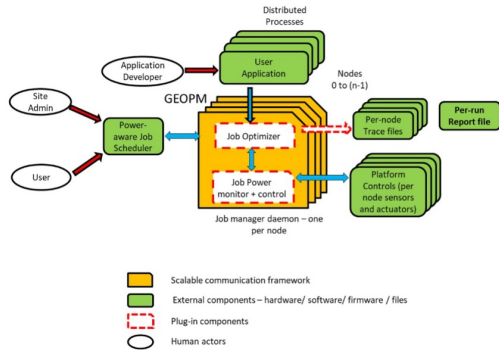


Fig. 3: High-Level Overview of GEOPM [1]

GEOPM [58], [74] is a community-driven, cross-platform, open source, job-level power management framework. It provides multiple interfaces to enable interoperability with external HPC software components such as enabling job schedulers and resource managers to drive job-aware system-wide power efficiency improvements in Figure 3. GEOPM enables control and monitoring of hardware/software knobs across multiple platforms and architectures such as leveraging multiple power and performance knobs like Intel’s hardware power-limiting capability (RAPL [77]) for achieved CPU frequency and instructions retired. Because the latest version of GEOPM (1.x) is installed on Theta but is not available on Summit because of special privilege requirement to access the low-level msr (model specific registers) counters and the power measurement of Power9 is not available to the public on Summit, Figure 4 shows the proposed framework for autotuning energy and EDP of various hybrid MPI/OpenMP applications on Theta. The average node energy consumed by the application is the primary metric.




Fig. 4: Framework for Autotuning Energy

This energy autotuning framework is similar to the performance framework with the five steps in Figure 1. Steps 1 and 2 are the same. There are some differences in Steps 3, 4, and 5. The GEOPM job launch script, `geopmlaunch` [74], queries and uses the `OMP_NUM_THREADS` environment variable to choose affinity masks for each process. The principal job of `geopmlaunch` to `aprun` is to set explicit per-process CPU affinity masks that will optimize performance while enabling the GEOPM controller thread to run on a core isolated from the cores used by the primary application. The `geopmlaunch` enables the GEOPM library to interpose on MPI using the PMPI interface through the `LD_PRELOAD` mechanism for unmodified binaries.

The iterative phase of the proposed autotuning framework in energy has the following steps:

- Steps Steps 1 and 2 are the same as shown in Figure 1.
- Step3 `ytopt` sets the `OMP_NUM_THREADS` environment variable, and generates the `aprun` command line for application launch.
- Step4 The dynamic linking is required with the `-dynamic` flag for the compiling.
- Step5 `ytopt` uses the `geopmlaunch` to launch the `aprun` command line with the options `--geopm-ctl=pthread,` which launches the controller as an extra `pthread` per node, and `--geopm-report=gm.report,` which creates the summary

report file gm.report about performance, power, and energy for each node to evaluate the application with the configuration. ytopt processes the summary report file from GEOPM to record the average node energy in the performance database.

Steps 1–5 are repeated until the maximum number of code evaluations or the wall-clock time for the run on Theta.

V. AUTOTUNING MIXED PRAGMAS ON A SINGLE NODE

In this section we apply the proposed framework in Figure 1 to autotune XSBench with mixed pragmas on a single node of Theta and Summit. We use its OpenMP version on Theta and its OpenMP offload version on Summit.

A. On Theta

We modify the OpenMP version of XSBench by adding more OpenMP pragmas and Clang loop optimization pragmas, such as loop unrolling and tiling [53]. We integrate the new OpenMP pragmas with Clang loop pragmas as parameters to autotune the XSBench and to make sure that the result is verified. Note that we use the clang-14 compiler from SOLLVE LLVM [56] to compile the original and the mixed XSBench on Theta.

We use 9 parameters to define the following parameter space. The system runtime parameters are OMP_NUM_THREADS, OMP_PLACES, OMP_PROC_BIND, and OMP_SCHEDULE; the unique application parameters are block size for OpenMP dynamic schedule, Clang loop unrolling full, additional OpenMP parallel for, and two tile sizes for 2D loop tiling for a double nested loop (because this loop fails when parallelizing it in OpenMP). Because each Theta node has 64 cores with up to 4 threads per core, we choose 10 choices for OMP_NUM_THREADS in the range of 4 to 256 threads. The OpenMP specification includes many environment variables related to program execution [66]. For the OMP_PLACES, there are three options: cores (threads are allowed to float on cores), threads (threads are bound to specific logical processors), and sockets (threads are allowed to float on sockets). For the OMP_PROC_BIND, there are also three options: close (threads placed consecutively), spread (threads spread equally on hardware), and master (threads placed on master to enhance locality). OMP_SCHEDULE allows specifying the schedule type (static, dynamic, or auto) with the default chunk size. For the block size (default 100 from the original code), we choose 12 choices in the range from 10 to 400. For the unrolling and additional OpenMP parallel for (4 in total), each has two choices with or without it. For two tile sizes for 2D loop tiling, we choose 11 choices in the range from 2 to 1,024 for each dimension. Therefore, the parameter space with total different configurations is $270 \times 5808 \times 4 = 6,272,640$, as shown in Table III.

Figure 5a shows the autotuning of the mixed pragmas version of XSBench (history based) on a single node, where wall-clock time stands for the time from the start of the autotuning to its end; the red line stands for the baseline runtime (3.31 s for the original code using 64 threads); and

the blue line stands for the autotuning process over time. We achieve the best performance 3.262 s, and the search reaches the good region of the parameter space over time. Figure 5c shows the ytopt overhead for each evaluation during the autotuning. The overhead is less than 65 s for the large parameter space.

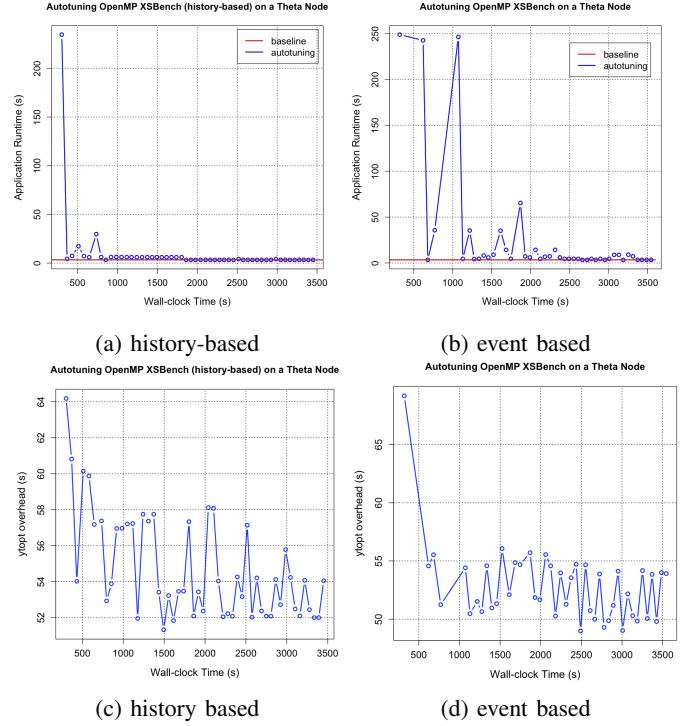


Fig. 5: Autotuning the Mixed-Pragmas Version of XSBench on a Theta Node

Figure 5b shows autotuning the mixed-pragmas version of XSBench (event based) on a single node. We observe the best performance 3.339 s (the baseline: 3.395 s for using 64 threads), and the search reaches the good region of the parameter space over time. Figure 5d shows the ytopt overhead for each evaluation over time. The overhead is between 49 s and 69.2 s for the large parameter space. We observe that the ytopt overhead for the first evaluation is the largest because it also includes setting the ytopt conda environment. Overall, the ytopt overhead is less than 70 s.

B. On Summit

We use the OpenMP offload version of XSBench to autotune the application on a Summit node. We integrate the additional OpenMP pragmas with some clauses as parameters to autotune the XSBench and to make sure that the result is verified. The OpenMP offload version supports only the event-based simulation. In the rest of this paper, we use XSBench with the event-based simulation for our experiments.

We use 9 parameters to define the following parameter space. The system runtime parameters are OMP_NUM_THREADS, OMP_PLACES, OMP_PROC_BIND, OMP_SCHEDULE, and OMP_TARGET_OFFLOAD;

the unique application parameters are additional “#pragma omp parallel for”, simd, device, and schedule clauses. Because each Summit node has 42 cores with up to 4 threads per core, we choose 10 choices for OMP_NUM_THREADS in the range from 4 to 168 threads. OMP_TARGET_OFFLOAD affects the behavior of execution on host and device including host fallback and provides three options: DEFAULT (try to execute on a GPU; if a supported GPU is not available, fall back to the host), DISABLED (do not execute on the GPU even if one is available; execute on the host), and MANDATORY (execute on a GPU or terminate the program). The simd clause is to create a team of threads to execute the loop in parallel using SIMD instructions. The device clause is to evaluate an assigned non-negative integer value less than the value of omp_get_num_devices() (6 devices on a Summit node). schedule(static,1) for the OpenMP target pragmas is for memory access coalescing; scheduling a chunk size of 1 for each thread allows consecutive threads to access consecutive global memory locations. We choose six chunk sizes in the range from 1 to 32, adding one of them or without adding one as the total 7 choices for the parameter schedule. Therefore, the parameter space with all the different configurations is $810 \times 56 \times 4 = 181,440$, as shown in Table III.

Figure 6 shows the autotuning of the OpenMP offload version of XSBench (event based) on a single Summit node using 6 GPUs. We observe the best performance 2.138 s (the baseline: 2.20 s for using 6 GPUs and 168 threads), and the search reaches the good region of the parameter space over time in Figure 6a. Figure 6b shows the ytopt overhead for each evaluation during the autotuning. It is less than 24 s. This is much faster than on Theta.

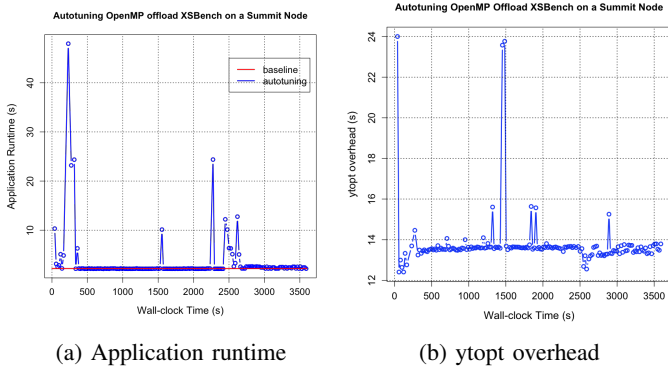


Fig. 6: Autotuning OpenMP Offload Version of XSBench (Event Based) on a Summit Node

VI. AUTOTUNING PERFORMANCE AT LARGE SCALES

In this section we apply the proposed framework in Figure 1 to autotune the performance of four ECP proxy applications—XSBench, AMG, SWFFT, and SW4lit—on both Theta and Summit. To launch an application to compute nodes, Theta uses aprun, and Summit uses jsrun. The processor core on both Theta and Summit supports the simultaneous multithreading (SMT) level of 4 as default so that the number of threads per node is supported up to 256 on Theta and up to 168 on

Summit. In Step 3 shown in Figure 1, based on the value of the number of threads from the selected configuration, the number of nodes reserved, and the number of MPI ranks, ytopt generates the aprun/jsrun command line for application launch on compute nodes. For instance, we reserve 4,096 nodes with one MPI rank per node to run an application on Theta and Summit.

On Theta we use the following algorithm to generate an aprun command line.

```
OMP_NUM_THREADS=n
if (n <= 64) {
  aprun -n 4096 -N 1 -cc depth -d n -j 1 application
} else { if (n <= 128) {
  aprun -n 4096 -N 1 -cc depth -d n/2 -j 2 application
} else { if (n <= 192) {
  aprun -n 4096 -N 1 -cc depth -d n/3 -j 3 application
} else {
  aprun -n 4096 -N 1 -cc depth -d n/4 -j 4 application
}
}
}
```

When we choose the number of threads n for each case, we make sure that $n/2$, $n/3$, or $n/4$ is integer on Theta. Then we use the algorithm to set the proper number of threads per core to generate the aprun command line.

On Summit we use the following algorithm to generate the jsrun command line. When the application uses 6 GPUs per node for the hybrid MPI/OpenMP offload application XSBench, we use the algorithm.

```
OMP_NUM_THREADS=n
jsrun -n4096 -a6 -g6 -c42 -bpacked:n/4 -dpacked application
```


When we choose the number of threads n , we make sure that $n/4$ is an integer because of the SMT level of 4 as default on Summit. We set one MPI rank per GPU and 42 cores per node for threads.

When the application uses only CPUs per node without any GPU for the hybrid MPI/OpenMP applications AMG, SWFFT, and SW4lit, we use the following algorithm to set one MPI rank per node and 42 cores per node for threads.

```
OMP_NUM_THREADS=n
jsrun -n4096 -a1 -g0 -c42 -bpacked:n/4 -dpacked application
```

For measuring the baseline performance for each application with a given problem size, we set the number of threads to 64 (which results in the best performance) on Theta and 168 threads (which also results in the best performance) on Summit to run the application under the default system configuration five times. Then we use the smallest application runtime as the baseline for the application. Notice that because of the limited node-hour allocations on Theta and Summit for our projects, we had to set most of the wall-clock times for autotuning runs at half an hour (1800 s). This limits the number of evaluations for different configurations during the autotuning.

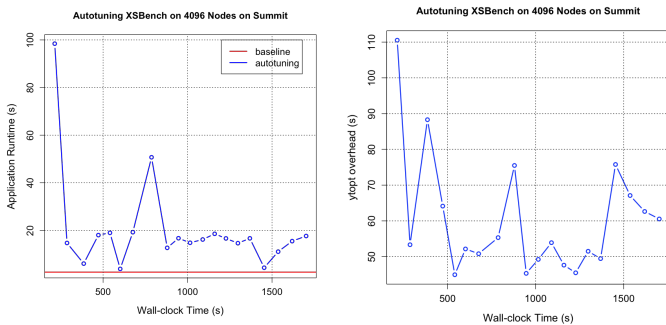
Figure 7 shows autotuning MPI/OpenMP XSBench with the large problem size on 1,024 and 4,096 nodes on Theta. Because XSBench is weak scaling, both the autotuning processes are similar because of its embarrassingly parallel implementation of XSBench. We observe that the ytopt search reaches the good region of the parameter space over time that is close to that of the baseline. The ytopt overhead is similar to that in Figure 5d on Theta.



(a) on 1024 nodes (b) on 4096 nodes

Fig. 7: Autotuning XSBench at Large Scales on Theta

Figure 8 shows autotuning MPI/OpenMP offload XSBench using 6 GPUs per node and 1 MPI rank per GPU on 4096 nodes on Summit. We observe that the ytopt search gradually reaches the good region of the parameter space over time (baseline: . Because of the limited number of evaluations (20), however, it does not reach the optimal performance yet in Figure 8a. Figure 8b shows the ytopt overhead during the entire autotuning. Notice that the first ytopt overhead (111 s) also includes the time spent in setting the ytopt conda environment and loading the nvhpc module, however, most of the times are around 60 s. The ytopt overhead is less than 111 s.



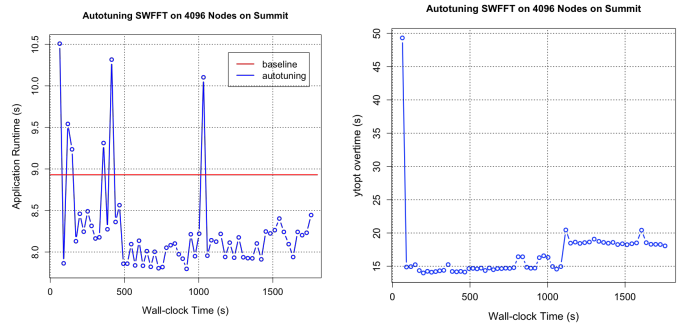
(a) on 4096 nodes (b) ytopt overhead

Fig. 8: Autotuning XSBench at Large Scale on Summit

SWFFT is weak scaling. Figure 9 shows autotuning SWFFT with a problem size of 3D grid 4096x4096x4096 on 4,096 nodes on Summit. We observe that the ytopt search reaches the good region of the parameter space over time with the smallest runtime of 7.797 s that is better than the baseline (8.93s) in Figure 9a. This is 12.69% performance improvement. The ytopt overhead is shown in Figure 9b, and most of the times are around 20 s because of the small parameter space for SWFFT. So the ytopt overhead is less than 50 s.


Figure 10 shows autotuning SWFFT with the same problem size on 4,096 nodes on Theta. We observe that the ytopt search reaches the good region of the parameter space over time that is close to the time of the baseline in Figure 10a. The ytopt overhead is less than 30 s in Figure 10b.

AMG is weak scaling. Figure 11 shows autotuning of AMG



(a) on 4096 nodes (b) ytopt overhead

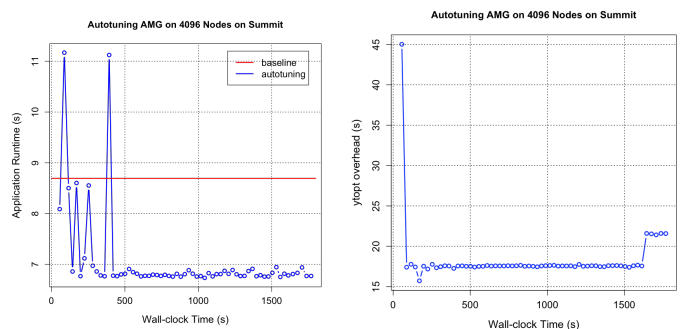
Fig. 9: Autotuning SWFFT at Large Scale on Summit



(a) on 4096 nodes (b) ytopt overhead

Fig. 10: Autotuning SWFFT at Large Scale on Theta

on 4,096 nodes on Summit. We use the 3D Laplace problem "laplace -n 100 100 100 -P 16 16 16" as the input, which means generating a problem with 1,000,000 grid points per MPI rank with a domain size 1600 x 1600 x 1600 on 4,096 nodes with 1 MPI rank per node and various numbers of threads per MPI rank. We observe that the ytopt autotuning reaches the best configuration with the smallest runtime of 6.734 s, which is much better than the baseline performance of 8.694 s in Figure 11a. This is a 22.54% performance improvement. Figure 11b shows the ytopt overhead is less than 45 s.



(a) on 4096 nodes (b) ytopt overhead

Fig. 11: Autotuning AMG at Large Scale on Summit

Figure 12 shows the autotuning of AMG on 4,096 nodes on Theta. Because of the limited wall-clock time (1800 s), we see only six evaluations on 4,096 nodes, mainly caused

by the second very long evaluation (1039.06 s) in Figure 12a. We find the configuration for the long evaluation includes system parameters: 48 threads; OMP_PLACES=threads (that are bound to specific logical processors); OMP_PROC_BIND=master (threads placed on master place to enhance locality); and OMP_SCHEDULE=dynamic. We observe that the system parameter setting mainly causes the long application runtime because the first 48 cores of 64 cores are used and every two cores share the L2 cache. Figure 12b still shows that the ytopt overhead is less than 34 s.

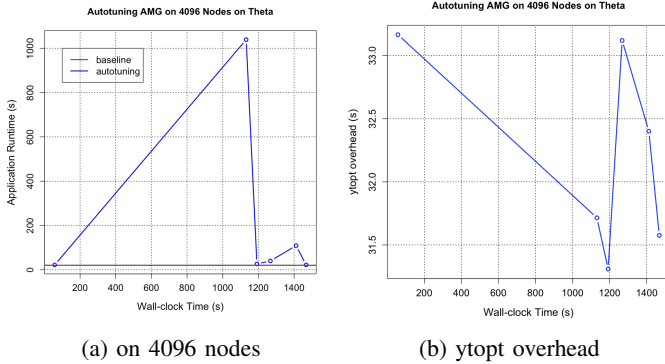


Fig. 12: Autotuning AMG at Large Scale on Theta

SW4lite with the large problem LOH.1-h50 is strong scaling so that we can test it on up to 1,024 nodes. Figure 13 shows autotuning SW4lite on 1,024 nodes on Summit. As described in Table III, the parameter space size for this application is 2,211,840. As shown in Figure 13a, at the beginning of the autotuning, ytopt samples the parameter space randomly for initial evaluations, then leverages the surrogate model to balance exploration of the search space and identifies more-promising parameter configurations using the LCB acquisition function. We observe that ytopt reaches the best configuration with the smallest runtime of 7.661 s, which is much better than the baseline performance of 11.067 s. This is a 30.78% performance improvement. Figure 13b shows the ytopt overhead is less than 46 s.




Fig. 13: Autotuning SW4lite at Large Scale on Summit

Figure 14a shows how SW4lite is autotuned on 1,024 nodes on Theta. We observe that ytopt reaches the best configuration with the smallest runtime of 14.427 s, which is much better

than the baseline performance of 171.595 s. This is a 91.59% performance improvement. We achieve the large improvement because we use the improved version of SW4lite [64] to define the parameter space for SW4lite with the parameter MPI_Barrier(MPI_COMM_WORLD). When running SW4lite on 1,024 nodes to measure the baseline performance, the compute time is small (around 3 s), but the communication time increases significantly (around 168 s) on Theta for the original code. Figure 14b shows the ytopt overhead during the entire autotuning which is less than 46 s.

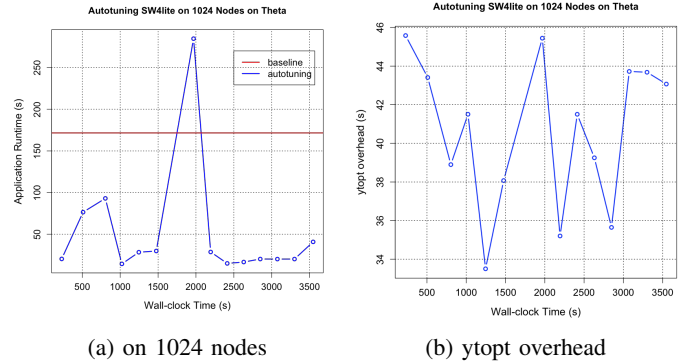


Fig. 14: Autotuning SW4lite at Large Scale on Theta

Overall, we observe that the ytopt overheads for the four applications are impacted mainly by the systems (for launching the application on the compute nodes) and application compiling times (given in Table II). We find that the ytopt overhead on up to 4,096 nodes on both Theta and Summit is less than 111 s in Table IV. This shows that our autotuning framework has low overhead and good scalability because the ytopt overhead does not increase much for autotuning the applications on small or large number of nodes.

TABLE IV: The maximum ytopt overhead (seconds) for each application on Theta and Summit

System	XSbench-Mixed	XSbench	SWFFT	AMG	SW4lite
Theta	70	69	30	34	46
Summit	24	111	50	45	46

VII. AUTOTUNING ENERGY AT LARGE SCALES

In this section we apply the proposed energy autotuning framework in Figure 4 to autotune the energy and EDP of four ECP proxy applications—XSbench, AMG, SWFFT, and SW4lite—on up to 4,096 nodes on Theta. Because energy consumption captures the tradeoff between the application runtime and power consumption and EDP captures the tradeoff between the application runtime and energy consumption, we use the autotuning framework to explore these tradeoffs for energy efficient application execution.

For measuring the baseline energy for each application with a given problem size, we set the number of threads to 64 on Theta and use GEOPM to run the application under the default system configuration five times. Then we use the smallest energy as the baseline for the application.

For autotuning energy or EDP, after the evaluation of a configuration GEOPM generates the summary report gm.report, which records the package energy and DRAM energy for each node; we accumulate these as the node energy. When ytop receives the report from GEOPM, it calculates an average node energy and uses that average energy as the primary metric for autotuning. Similarly, the average EDP is calculated.

Figure 15 shows using our energy framework to autotune the energy of the four ECP proxy applications at large scales on Theta. Figure 15a presents autotuning the energy of XSBench on 4,096 nodes, where the red line stands for the baseline node energy of 2494.905J. Using the framework achieves the lowest energy of 2280.806J. This is an 8.58% energy savings. Figure 15b shows autotuning the energy of SWFFT on 4,096 nodes. The baseline node energy for SWFFT is 3185.027J. Using the framework achieves the lowest node energy of 3118.604J. This is a 2.09% energy savings.

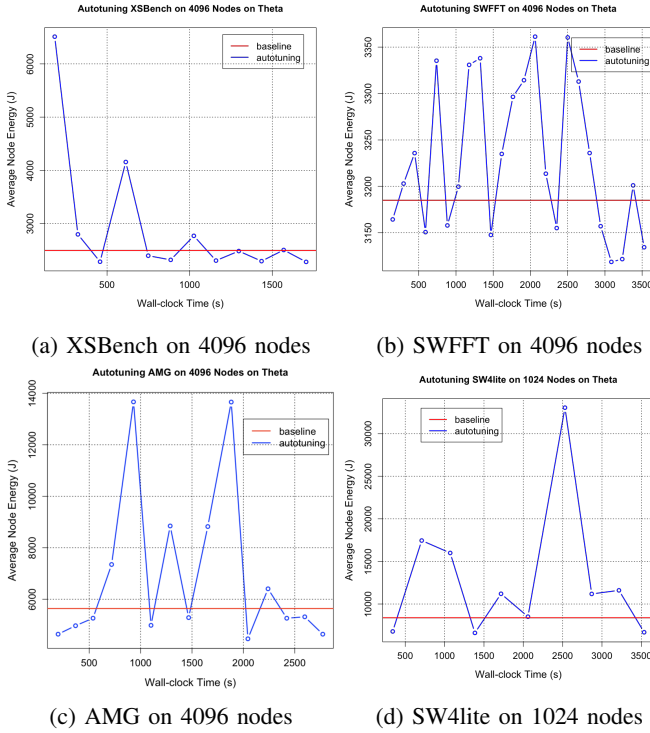


Fig. 15: Autotuning Energy at Large Scales on Theta

Figure 15c presents autotuning the energy of AMG on 4,096 nodes. The baseline node energy for AMG is 5642.568J. Using the framework achieves the lowest node energy of 4566.747J. This is a 20.88% energy saving. Figure 15d presents autotuning the energy of SW4lite on 1,024 nodes. The baseline node energy for SW4lite is 8384.034J. Using the framework achieves the lowest node energy of 6606.233J. This is a 21.20% energy saving. Compared with Figure 14a for SW4lite, we identified the best configuration (32, 'sockets', 'spread', 'static', ' ', ' ', '#pragma omp for nowait', ' ') which resulted in 91.59% performance improvement. The same configuration also resulted in the 21.20% energy saving because the large performance improvement led to the energy saving. As we

discussed before, the application runtime for SW4lite on 1024 nodes was dominated by the low power communication for the baseline, this was why the energy saving percentage is much less than the performance improvement percentage. Based on our observation, this is the case for other applications.

Figure 16 shows using our energy framework to autotune the EDP of the four ECP proxy applications at large scales on Theta. Figure 16a presents autotuning the energy of XSBench on 4,096 nodes, where the red line stands for the baseline node EDP. Using the framework achieves the lowest EDP with 37.84% improvement. Figure 16b shows autotuning the energy of SWFFT on 4,096 nodes. Using the framework achieves the lowest EDP with 5.24% improvement. Figure 16c presents autotuning the energy of AMG on 4,096 nodes. Using the framework achieves the lowest EDP with 24.13% improvement. Figure 16d presents autotuning the energy of SW4lite on 1,024 nodes. Using the framework achieves the lowest EDP with 23.70% improvement. Because EDP is the product of energy and application runtime, the EDP improvement is better than the energy improvement shown in Table V. The best configuration for using EDP as the metric is similar to that for using energy as metric.

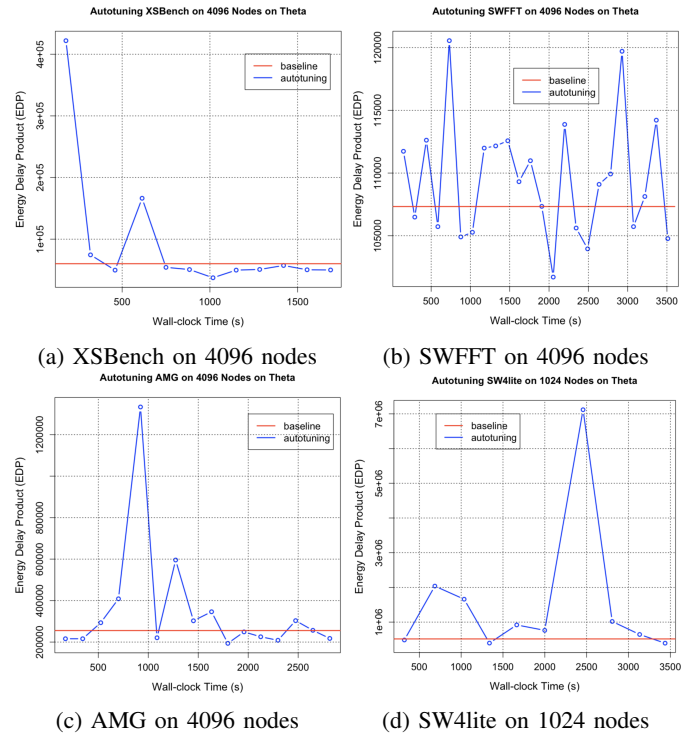


Fig. 16: Autotuning Energy Delay Product at Large Scales on Theta

Overall, using our energy autotuning framework to identify the best configurations for the four ECP proxy applications results in up to 21.2% energy savings and up to 37.84% improvement in EDP on up to 4,096 nodes shown in Table V. This aids us in exploring the tradeoffs between application runtime and power/energy for energy efficient application execution.

TABLE V: Improvement percentage (%) for each application on Theta

Theta	XSbench	SWFFT	AMG	SW4lite
Energy	8.58	2.09	20.88	21.20
EDP	37.84	5.24	24.13	23.70

VIII. CONCLUSIONS

In this paper, we proposed the low-overhead autotuning frameworks to autotune four hybrid MPI/OpenMP ECP proxy applications—XSbench, AMG, SWFFT, and SW4lite—at large scales and explored the tradeoffs between application runtime and power/energy for energy efficient application execution. We used Bayesian optimization with a Random Forest surrogate model to effectively search the parameter spaces with up to 6 million different configurations on Theta and Summit. We used the autotuning framework to explore the tradeoffs between application runtime and power/energy for energy efficient application execution. The experimental results showed that our autotuning framework had low overhead and good scalability. By using the autotuning framework to identify the best configuration, we achieved up to 91.59% performance improvement, up to 21.2% energy savings, and up to 37.84% EDP improvement on up to 4,096 nodes. The topt autotuning framework is open source and available to download from the link in [46].

For future work, we will improve the framework overhead by reducing the application compiling time with pre-compiling the unchanged code files and setting a proper evaluation timeout to evaluate more good configurations. Our current autotuning framework uses Ray [78] to do one evaluation each time; this affected the effectiveness of identifying the promising search regions at the beginning of the autotuning. We plan to extend the framework to do multiple evaluations in parallel using libensemble [79] to improve the initial effectiveness. We also plan to add transfer learning and online tuning to the framework so that it can transfer what it learns from the applications at a small scale in problem sizes and system sizes to guide and/or predict the best configurations for autotuning at large scales.

ACKNOWLEDGMENTS

This work was supported in part by DOE ECP PROTEAS-TUNE, in part by DOE ASCR RAPIDS2, and in part by NSF grant CCF-2119203. We acknowledge the Argonne Leadership Computing Facility (ALCF) for use of Cray XC40 Theta under ALCF projects EE-ECP and Intel, and the Oak Ridge Leadership Computing Facility for use of Summit under the projects CSC383, MED106 and AST136. We also acknowledge Adrian Pope at ALCF for providing the SWFFT problem sizes. This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357. Development of the GEOPM software package has been partially funded through contract B609815 with Argonne National Laboratory.

REFERENCES

- [1] A. Bartolini, S. Brink, D. Cesarini, D. Ellsworth, R. Grant, S. Jana, M. Kondo, E. K. Lee, M. Maiterth, A. Marathe, T. Patki, S. Perarnau, V. Reis, M. Schulz, O. Vysocky, T. Wilde, and X. Wu, “White paper on powerstack.” [Online]. Available: <https://hpcpowerstack.github.io/raitenhaslach20.html>
- [2] C. Cantalupo, J. Eastep, S. Jana, M. Kondo, M. Maiterth, A. Marathe, T. Patki, B. Rountree, R. Sakamoto, M. Schulz, and C. Trinitis, “A strawman for an HPC PowerStack,” 8 2018. [Online]. Available: <https://www.osti.gov/biblio/1466153>
- [3] X. Wu, A. Marathe, S. Jana, O. Vysocky, J. John, A. Bartolini, L. Riha, M. Gerndt, V. Taylor, and S. Bhalachandra, “Toward an end-to-end auto-tuning framework in HPC PowerStack,” in *Proceedings of Energy Efficient HPC State of Practice 2020 (EE HPC SOP 20)*. Washington, DC, USA: IEEE Computer Society, 2020.
- [4] “ECP Proxy Applications Suite,” <https://proxyapps.exascaleproject.org/ecp-proxy-apps-suite/>.
- [5] “XSbench Version 19, ECP Proxy Application,” <https://github.com/ANL-CESAR/XSBench> (Downloaded in Oct. 2021).
- [6] “SWFFT Version 1.0, ECP Proxy Application,” <https://git.cels.anl.gov/hacc/SWFFT> (Downloaded in Oct. 2021).
- [7] “AMG Version 1.2, ECP Proxy Application,” <https://github.com/LLNL/AMG> (Downloaded in Oct. 2021).
- [8] “SW4lite Version 1.1, ECP Proxy Application,” <https://github.com/geodynamics/sw4lite> (Downloaded in Oct. 2021).
- [9] “Cray XC40 Theta, Argonne National Laboratory.” [Online]. Available: <https://www.alcf.anl.gov/theta>
- [10] “Summit, Oak Ridge National Laboratory.” [Online]. Available: <https://www.olcf.ornl.gov/olcf-resources/compute-systems/summit/>
- [11] R. C. Whaley and J. J. Dongarra, “Automatically tuned linear algebra software,” in *Proceedings of the 1998 ACM/IEEE Conference on Supercomputing*, ser. SC’98. Washington, DC, USA: IEEE Computer Society, 1998, pp. 1–27. [Online]. Available: <http://dl.acm.org/citation.cfm?id=509058.509096>
- [12] C. Tapus, I. Chung, and J. K. Hollingsworth, “Active harmony: Towards automated performance tuning,” in *SC ’02: Proceedings of the 2002 ACM/IEEE Conference on Supercomputing*, ser. ICS’02, 2002, p. 44.
- [13] I. Chung and J. K. Hollingsworth, “A case study using automatic performance tuning for large-scale scientific programs,” in *15th IEEE International Symposium on High Performance Distributed Computing (HPDC’06)*, ser. HPDC’06, 2006.
- [14] M. Gerndt and M. Ott., “Automatic performance analysis with Periscope,” *Concurrency and Computation: Practice and Experience*, vol. 22, no. 6, 2010.
- [15] C. Chen, J. Chame, and M. Hall, “Combining models and guided empirical search to optimize for multiple levels of the memory hierarchy,” in *International Symposium on Code Generation and Optimization*, March 2005.
- [16] A. Tiwari and J. K. Hollingsworth, “Online adaptive code generation and tuning,” in *2011 IEEE International Parallel & Distributed Processing Symposium (IPDPS’11)*, ser. IPDPS’11, 2011.
- [17] P. Balaprakash, R. B. Gramacy, and S. M. Wild, “Active-learning-based surrogate models for empirical performance running,” in *2013 IEEE International Conference on Cluster Computing (CLUSTER’13)*, ser. CLUSTER’13, 2013.
- [18] T. L. Falch and A. C. Elster, “Machine learning-based auto-tuning for enhanced performance portability of OpenCL applications,” *Concurrency and Computation: Practice and Experience*, vol. 29, no. 8, 2017.
- [19] A. H. Ashouri, W. Killian, J. Cavazos, G. Palermo, and C. Silvano, “A Survey on Compiler Autotuning Using Machine Learning,” *Computing Surveys (CSUR)*, vol. 51, no. 5, 2018.
- [20] W. F. Ogilvie, P. Petoumenos, Z. Wang, and H. Leather, “Minimizing the cost of iterative compilation with active learning,” in *2017 International Symposium on Code Generation and Optimization*, 2017.
- [21] A. Tiwari, C. Chen, J. Chame, M. Hall, and J. K. Hollingsworth, “A scalable auto-tuning framework for compiler optimization,” in *Proceedings of the 23rd IEEE International Parallel And Distributed Computing Symposium (IPDPS’09)*, ser. IPDPS’09. IEEE, 2009.
- [22] S. Muralidharan, M. Shantharam, M. Hall, M. Garland, and B. Catanzaro, “Nitro: A framework for adaptive code variant tuning,” in *2014 IEEE 28th International Parallel and Distributed Processing Symposium (IPDPS’14)*, ser. IPDPS’14, 2014.

- [23] V. Sreenivasan, R. Javali, M. Hall, P. Balaprakash, T. R. W. Scogland, and B. R. de Supinski, "A Framework for Enabling OpenMP Autotuning," in *OpenMP: Conquering the Full Hardware Spectrum (IWOMP'19)*, ser. IWOMP'19. Springer, 2019, pp. 50–60.
- [24] J. Katarzynski and M. Cytoński, "Towards autotuning of openmp applications on multicore architectures," *CoRR*, vol. abs/1401.4063, 2014. [Online]. Available: <http://arxiv.org/abs/1401.4063>
- [25] D. Mustafa, A. Aurangzeb, and R. Eigenmann, "Performance analysis and tuning of automatically parallelized OpenMP applications," in *Proceedings of the 7th International Conference on OpenMP in the Petascale Era*, ser. IWOMP'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 151–164. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2023025.2023041>
- [26] C. Liao, D. J. Quinlan, R. Vuduc, and T. Panas, "Effective source-to-source outlining to support whole program empirical optimization," in *Languages and Compilers for Parallel Computing (LCPC 2009)*, LNCS 5898, ser. LCPC2009, G. R. Gao, L. L. Pollock, J. Cavazos, and X. Li, Eds. Springer Berlin Heidelberg, 2010.
- [27] X. Wu, M. Kruse, P. Balaprakash, H. Finkel, P. Hovland, V. Taylor, and M. Hall, "Autotuning PolyBench Benchmarks with LLVM Clang/Polly loop optimization pragmas using Bayesian optimization," in *Proceedings of SC20 Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems*, ser. PMBS'20. Washington, DC, USA: IEEE Computer Society, 2020.
- [28] —, "Autotuning PolyBench Benchmarks with LLVM Clang/Polly loop optimization pragmas using Bayesian optimization," *Concurrency and Computation: Practice and Experience*, vol. Volume 34, Issue 20, e6683, <https://doi.org/10.1002/cpe.6683>, 2022.
- [29] Y. Liu, W. M. Sid-Lakhdar, O. Marques, X. Zhu, C. Meng, J. W. Demmel, and X. S. Li, "GPTune: multitask learning for autotuning exascale applications," in *Proceedings of PPOPP '21: Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP'21. New York, NY, USA: Association for Computing Machinery, February 2021, pp. 234–246.
- [30] R. B. Roy, T. Patel, V. Gadepally, and D. Tiwari, "Bliss: Auto-tuning complex applications using a pool of diverse lightweight learning models," in *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, ser. PLDI'21. New York, NY, USA: Association for Computing Machinery, June 20–25, 2021, pp. 1280–1295.
- [31] T. Miyazaki, I. Sato, and N. Shimizu, "Bayesian optimization of HPC systems for energy efficiency," in *High Performance Computing*, R. Yokota, M. Weiland, D. Keyes, and C. Trinitis, Eds. Cham: Springer International Publishing, 2018, pp. 44–62.
- [32] F.-J. Willemsen, R. van Nieuwpoort, and B. van Werkhoven, "Bayesian optimization for auto-tuning GPU kernels," in *2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, 2021, pp. 106–117.
- [33] A. Roy, P. Balaprakash, P. D. Hovland, and S. M. Wild, "Exploiting performance portability in search algorithms for autotuning," in *2016 IEEE International Parallel and Distributed Processing Symposium Workshops: iWAPT*, ser. iWAPT'16, 2016.
- [34] A. Marathe, R. Anirudh, N. Jain, A. Bhatle, J. Thiagarajan, B. Kaikhura, J. Yeom, B. Rountree, and T. Gambin, "Performance modeling under resource constraints using deep transfer learning," in *SC17*, ser. SC17, Nov. 2017.
- [35] J. Thiagarajan, N. Jain, R. Anirudh, A. Gimenez, R. Sridhar, A. Marathe, T. Wang, M. Emani, A. Bhatle, and T. Gambin, "Bootstrapping parameter space exploration for fast tuning," in *ICS'18*, ser. ICS'18, June 2018.
- [36] P. Bruel, S. Quito Masnada, B. Videau, A. Legrand, J.-M. Vincent, and A. Goldman, "Autotuning under Tight Budget Constraints: A Transparent Design of Experiments Approach," in *International Symposium in Cluster, Cloud, and Grid Computing (CCGrid'19)*, ser. CCGrid'19, May 2019.
- [37] A. Hartono, B. Norris, and P. Sadayappan, "Annotation-based empirical performance tuning using Orío," in *2009 IEEE International Symposium on Parallel Distributed Processing*, ser. ISDPD'09, 2009.
- [38] J. Ansel, S. Kamil, K. Veeramachaneni, J. Ragan-Kelley, J. Bosboom, U.-M. O'Reilly, and S. Amarasinghe, "OpenTuner: An extensible framework for program autotuning," in *Proceedings of the 23rd International Conference on Parallel Architectures and Compilation Techniques*, ser. PACT'14. ACM, Aug. 2014, pp. 303–316.
- [39] T. Nelson, A. Rivera, P. Balaprakash, M. Hall, P. D. Hovland, E. Jessup, and B. Norris, "Generating efficient tensor contractions for GPUs," in *2015 44th International Conference on Parallel Processing*, ser. ICPP'15, Sep. 2015.
- [40] G. P. E. T. Zingales, "HalideTuner: Generating and tuning Halide schedules with OpenTuner," Master's Thesis, Massachusetts Institute of Technology, 2015.
- [41] A. Rasch, M. Haidl, and S. Gortlach, "ATF: A generic auto-tuning framework," in *19th International Conference on High Performance Computing and Communications; 15th International Conference on Smart City; 3rd International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, ser. HPCC'17. IEEE, 2017, pp. 64–71.
- [42] C. Nugteren and V. Codreanu, "ClTune: A generic auto-tuner for opencl kernels," *CoRR*, vol. abs/1703.06503, 2017. [Online]. Available: <http://arxiv.org/abs/1703.06503>
- [43] P. Pfafe, T. Grosser, and M. Tillmann, "Efficient hierarchical online-autotuning: A case study on polyhedral accelerator mapping," in *Proceedings of the ACM International Conference on Supercomputing*, ser. ICS '19. New York, NY, USA: Association for Computing Machinery, 2019.
- [44] B. van Werkhoven, "Kernel tuner: A search-optimizing GPU code auto-tuner," *Future Generation Computer Systems*, vol. 90, 2019.
- [45] A. Haj-Ali, H. Genc, Q. Huang, W. Moses, J. Wawrzyniek, K. Asanovi?, and I. Stoica, "ProTuner: Tuning programs with Monte Carlo tree search," *arXiv preprint*, 2020. [Online]. Available: <http://arxiv.org/abs/2005.13685>
- [46] "ytopt," <https://github.com/ytopt-team/ytopt>, <https://github.com/ytopt-team/autotune/blob/master/Benchmarks>.
- [47] M. Kruse, H. Finkel, and X. Wu, "Autotuning search space for loop transformations," in *Proceedings of SC20 Workshop on the LLVM Compiler Infrastructure in HPC*. Washington, DC, USA: IEEE Computer Society, 2020.
- [48] J. Koo, P. Balaprakash, M. Kruse, X. Wu, P. Hovland, and M. Hall, "Customized Monte Carlo tree search for LLVM/Polly's composable loop optimization transformations," in *2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, 2021, pp. 82–93.
- [49] A. Rasch, R. Schulze, M. Steuwer, and S. Gortlach, "Efficient autotuning of parallel programs with interdependent tuning parameters via auto-tuning framework (ATF)," *ACM Trans. Archit. Code Optim.*, vol. 17, no. 5, 2021.
- [50] A. Adams, K. Ma, L. Anderson, R. Baghdadi, T.-M. Li, M. Gharbi, B. Steiner, S. Johnson, K. Fatahalian, F. Durand, and J. Ragan-Kelley, "Learning to Optimize Halide with Tree Search and Random Programs," *ACM Transactions on Graphics*, vol. 38, no. 4, Jul. 2019.
- [51] U. Beaugnon, B. Clement, N. Tollenare, and A. Cohen, "On the representation of partially specified implementations and its application to the optimization of linear algebra kernels on GPU," *arXiv preprint*, Apr. 2019, arXiv:1904.03383 [cs.PL].
- [52] "mctree," <https://github.com/Meinersbur/mctree>.
- [53] M. Kruse and H. Finkel, "User-directed loop-transformations in Clang," in *Proceedings of the 5th Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC 2018. November 12. Dallas, USA)*, ser. LLVM-HPC'18. IEEE, 2018.
- [54] "Clang: A C Language Family Frontend for LLVM," <http://clang.llvm.org>.
- [55] T. Grosser, A. Gröbflinger, and C. Lengauer, "Polly – Performing Polyhedral Optimizations on a Low-Level Intermediate Representation," *Parallel Processing Letters*, vol. 22, no. 04, 2012.
- [56] "SOLLVE LLVM," <https://github.com/SOLLVE/llvm-project.git>.
- [57] T. Yuki and L. Pouchet, "Polybench 4.2," <http://sourceforge.net/projects/polybench/>, May 9, 2016.
- [58] J. Eastep, S. Sylvester, C. Cantalupo, B. Geltz, F. Ardanaz, A. Al-Rawi, K. Livingston, F. Keceli, M. Maiterth, and S. Jana, "Global extensible open power manager: a vehicle for HPC community collaboration on co-designed energy management solutions," in *International Supercomputing Conference: High Performance Computing (ISC 2017)*, ser. LNCS, vol. 10266, 2017.
- [59] "NVIDIA system management interface (nvidia-smi)." [Online]. Available: <https://developer.nvidia.com/nvidia-system-management-interface>, 201
- [60] "OpenMC Monte Carlo Particle Transport Code," <https://github.com/openmc-dev/openmc>.

- [61] B. Sjogreen and N. Petersson, "A fourth order accurate finite difference scheme for the elastic wave equation in second order formulation," *Journal of Scientific Computing*, vol. 52, no. 1, pp. 17–48, 2012.
- [62] B. Sjogreen, "SW4 final report for iCOE," Technical Report LLNL-TR-759417, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, October 3, 2018.
- [63] S. Day and et al., "Tests of 3D elastodynamic codes, final report for lifelines project 1A01, Pacific Earthquake Engineering Center," 2001.
- [64] X. Wu, V. Taylor, and Z. Lan, "Performance and energy improvement of ECP proxy app SW4lite under various workloads," in *2021 IEEE/ACM Workshop on Memory Centric High Performance Computing (MCHPC)*. Los Alamitos, CA, USA: IEEE Computer Society, nov 2021, pp. 17–24. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/MCHPC54807.2021.00009>
- [65] M. Lindauer, K. Eggensperger, M. Feurer, A. Biedenkapp, J. Marben, P. Müller, and F. Hutter, "Boah: A tool suite for multi-fidelity Bayesian optimization & analysis of hyperparameters," *arXiv:1908.06756 [cs.LG]*, 2019. [Online]. Available: <https://automl.github.io/ConfigSpace/master/>
- [66] "Openmp application programming interface version 5.0," November 2018. [Online]. Available: <https://www.openmp.org/wp-content/uploads/OpenMP-API-Specification-5.0.pdf>
- [67] "Hpc powerstack." [Online]. Available: <https://hpcpowerstack.github.io>, <https://powerstack.caps.in.tum.de>
- [68] R. E. Grant, M. Levenhagen, S. L. Olivier, D. DeBonis, K. T. Pedretti, and J. H. L. III, "Standardizing power monitoring and control at exascale," *IEEE Computer*, vol. 49, no. 10, pp. 38–46, Oct. 2016.
- [69] R. E. Grant, B. Rountree, J. Hansen, and et al., "High performance computing power application programming interface specification community version 1.0, technical report." [Online]. Available: https://github.com/pwrapi/powerapi_spec/releases/
- [70] "Intelligent platform management interface (IPMI), ipmi technical resources." [Online]. Available: <https://www.intel.com/content/www/us/en/servers/ipmi/ipmi-technical-resources.html>
- [71] "Distributed management task force, inc. (DMTF), Redfish white paper." [Online]. Available: https://www.dmtf.org/sites/default/files/standards/documents/DSP2044_1.0.4.pdf
- [72] M. Maiterth, G. Koenig, K. Pedretti, S. Jana, N. Bates, A. Borghesi, D. Montoya, A. Bartolini, and M. Puzovic, "Energy and power aware job scheduling and resource management: Global survey — initial analysis," in *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2018, pp. 685–693.
- [73] "Energy efficient hpc working group." [Online]. Available: <https://eehpcwg.llnl.gov/>
- [74] "GEOPM: Global extensible open power manager." [Online]. Available: <https://geopm.github.io>, <https://github.com/geopm/theta-tutorial/blob/master/QUICKSTART.md>
- [75] H. David, E. Gorbato, U. R. Hanebutte, R. Khanna, and C. Le, "RAPL: Memory power estimation and capping," in *2010 ACM/IEEE International Symposium on Low-Power Electronics and Design (ISLPED)*, 2010, pp. 189–194.
- [76] X. Wu, A. Marathe, and S. Jana, "End-to-End PowerStack Codesign for Energy Efficient HPC (position paper)," 2021 DoE Advanced Scientific Computing Research (ASCR) Workshop on Reimagining Codesign (ReCoDe), March 16-18, 2021.
- [77] Intel Corporation, "Intel(R) 64 and IA-32 Architectures Software Developer's Manual Vol. 3B: System Programming Guide, Part-2," February 2014. [Online]. Available: <http://www.intel.com/content/www/us/en/architecture-and-technology/64-ia-32-architectures-software-developer-vol-3b-part-2-manual.html>
- [78] "Ray." [Online]. Available: <https://github.com/ray-project/ray>, <https://docs.ray.io/en/latest/>
- [79] "libEnsemble: a Python library to coordinate the concurrent evaluation of dynamic ensembles of calculations." [Online]. Available: <https://github.com/Libensemble/libensemble>, <https://libensemble.readthedocs.io/en/main/>