Lawrence Berkeley National Laboratory

Recent Work

Title

ExaSAT: An exascale co-design tool for performance modeling

Permalink

<https://escholarship.org/uc/item/12r566sn>

Journal

International Journal of High Performance Computing Applications, 29(2)

ISSN

1094-3420

Authors

Unat, D
Chan, C
Zhang, W
et al.

Publication Date

2015-05-04

DOI

10.1177/1094342014568690

Peer reviewed

ExaSAT: An Exascale Co-Design Tool for Performance Modeling

Didem Unat, Cy Chan, Weiqun Zhang, Samuel Williams, John Bachan,

John Bell, and John Shalf

Lawrence Berkeley National Laboratory

dunat, cychan, weiqunzhang, swwilliams, jdbachan, jbbell, jshalf @lbl.gov

Abstract

One of the emerging challenges to design HPC systems is to understand and project the requirements of exascale applications. In order to determine the performance consequences of different hardware designs, analytic models are essential because they can provide fast feedback to the co-design centers and chip designers without costly simulations. However, current attempts to analytically model program performance typically rely on the user manually specifying a performance model. We introduce the ExaSAT framework that automates the extraction of parameterized performance models directly from source code using compiler analysis. The parameterized analytic model enables quantitative evaluation of a broad range of hardware design trade-offs and software optimizations on a variety of different performance metrics, with a primary focus on data movement as a metric. We demonstrate the ExaSAT framework's ability to perform deep code analysis of a proxy application from the DOE Combustion Co-design Center to illustrate its value to the exascale co-design process. ExaSAT analysis provides insights in the hardware and software tradeoffs and lays the groundwork for exploring a more targeted set of design points using cycle-accurate architectural simulators.

1 Introduction

The design of exascale systems are faced with challenges introduced by system cost and power consumption (Shalf et al., 2010). In order to improve delivered performance for large-scale applications within practical cost and power budgets, it is essential to move towards a hardware/software *co-design* process where the hardware design space is explored in tandem with software optimizations. The US Department of Energy co-design centers (Cesar, 2013; Exact, 2013; ExMatEx, 2013) are performing multi-disciplinary research to iteratively design various aspects of applications including core algorithms, programming models, compilers, and runtimes to ensure that they will meet the requirements of future scientific simulations. Co-design requires a performance framework to *rapidly* evaluate the proposed hardware and software changes and provide end-to-end analysis of an application.

In order to evaluate the hardware-software design trade-offs, we introduce a compiler-based performance modeling framework, ExaSAT (**Ex**ascale **S**tatic **A**nalysis **T**ool), that enables rapid exploration of hardware design space and helps bridge the communication gap between the application developers and hardware designers. Because many exascale architectural specifications are currently undefined, our performance model is parameterized to help explore different design choices. Additionally, our framework explores a parameterized software optimization space (e.g. cache blocking, fusion, etc.) together with the hardware design space so that we do not base conclusions about hardware requirements on unoptimized codes.

The ExaSAT framework focuses on structured problems from combustion codes. Combustion currently provides 85% of the nation’s energy needs and is a key driver for exascale computing (U.S. Energy Flow Trends, 2012). Economic and environmental concerns are driving the development of new combustion systems targeted toward clean and efficient use of alternative fuels. Developing these systems requires simulations with sufficient chemical fidelity to differentiate behavior of candidate fuels in realistic engine conditions. Exascale computing offers the promise of enabling the underlying science to design fuel efficient, clean burning vehicles, planes, and power plants for electricity generation (Exascale Research PI Meeting (2012)). For example, exascale computing will enable the development of new Homogenous Charge Compression Ignition (HCCI) engine designs that lead to lower emissions, cleaner combustion and 25-50% increase in efficiency. It is predicted that HCCI will require 20 days

runtime at billion way concurrency, 3 PB memory to hold the simulation state, and will generate 1.0 EB of data for analysis. Thus, studying performance requirements of combustion applications on potential exascale designs are extremely valuable.

By restricting the framework to structured grid problems, we improve the accuracy of the performance model. For example, while the compiler front-end gathers the read/write properties of streaming arrays, the cache model takes into account the reuse in stencil arrays when estimating the memory traffic. Since the access pattern of such applications operating on dense arrays can be statically inferred, we can quickly derive fast analytic performance models. Our approach does not support analysis for irregular or graph-based codes where the access pattern is only available at runtime. It is important to note that ExaSAT is not intended to replace architectural simulations. Rather it is intended to be used in conjunction with those other tools to prune the search space. Our framework provides a *missing capability* in the co-design toolset where fast evaluation is needed at the expense of accuracy. Simulations are slow, leading to very narrow, yet highly detailed analysis of a small kernel or a sub-component of a system. For example, it would take a hardware simulator (e.g. RAMP (Krasnov et al., 2007; Wawrzynek et al., 2007)) a few hours to generate a single configuration of a multicore processor, though the application on the configured architecture would then run in real-time. It is easier to configure a cycle-accurate software simulator (e.g. gem5 (Binkert et al., 2011)), but it would take several hours to run an application to get meaningful results. In comparison, ExaSAT can evaluate hundreds of hardware/software configurations per minute on a desktop machine. Thus, ExaSAT complements hardware and software simulators in the co-design process by serving as a design space pruning tool.

This paper makes the following contributions:

- We developed the *ExaSAT* framework to statically analyze an application and automatically gather key characteristics about the computation, communication, data access patterns and data locality that are important in characterizing performance of combustion codes.
- We designed an XML internal representation to represent the application workload and an XML machine specification to represent the exascale machine configuration. The XML serves as a medium and an interface for our framework to work with other tools, such as PIN tools or architectural simulators.

- We implemented a performance model that can combine both hardware and software parameters to bound performance, rapidly explore design trade-offs, and extrapolate these requirements to potential hardware realizations in the exascale timeframe (2020).
- We performed deep code analysis of SMC, a proxy implementation of a production combustion code, and used our results to address key co-design questions acquired from our industry partners. Finally, we quantified the SMC performance on exascale proxy architectures using ExaSAT.

The rest of the paper is organized as follows. Section 2 provides background on related work and explains how ExaSAT differs from existing performance modeling tools. Section 3 introduces the ExaSAT framework including the compiler-based front-end, XML specification for the code description and abstract machine model, and performance modeling component. Validation of the framework is provided at the end of the section. Section 4 provides an overview on the characteristics of combustion applications and gives details about a proxy application used to conduct performance analysis for this paper. We present performance analysis and results in Section 5. Section 6 makes projections on an exascale machine, evaluates the implications of our findings, and provides feedback to hardware/software designers for exascale systems. The section includes discussion on limitations and future work. We conclude the paper in Section 7.

2 Related Work

The overarching goal of the co-design centers is to understand the interplay between hardware and software design trade-offs. Given the uncertainty in exascale architecture, co-design centers need an application characterization tool to iteratively perform a hardware/software optimization process envisioned for the co-design of HPC systems. GEM5 (Binkert et al., 2011), CACTI (Thoziyoor et al., 2008), or SST (Rodrigues et al., 2011) are software simulators that parameterize the machine specifications but they are slow, leading to narrow analysis of small kernels or isolated components of the system such as the interconnect. FPGA-based cycle-accurate, circuit-level emulators such as RAMP (Krasnov et al., 2007; Wawrzyniek et al., 2007) and the CoDEx emulator (Shalf et al., 2011) can capture very detailed behavior of the architecture, but are not as easily configurable as software simulators. For

example, if the number of cores in the emulator is changed from 64 to 128, every single module will need to be manually adjusted for the new cache sizes, address spaces, and network sizes. Furthermore, very fine grained circuit level design introduces the danger of missing general performance trends because of the extraneous amounts of data generated. Benchmarking provides immediate response, but limits the analysis to current hardware architectures and the results can be biased toward particular implementation or compiler options used because we cannot separate implementation-specific results from performance opportunities.

Given the cost of setting up both simulations and emulations, analytic models play a complementary role in design space exploration to identify the subset that is of interest for further study with simulation and emulation. Higher level analytic models such as the Roofline model (Williams et al., 2009) provide a speed-of-light (cannot-exceed) performance expectations, but offer a very coarse-grained description of performance in terms of flops rates and DRAM bandwidth. Convolution-based approaches such as PMAC tools (Snaveley et al., 2002; Carrington et al., 2003) provide coarse-grained performance analysis through correlation and generate models by convolving application characteristics (the “signature”) through instrumentation with a vector describing the target machine attributes. Similar to ExaSAT, Pbound (Narayanan et al., 2010) mixes static and runtime data to estimate upper performance bounds. However, ExaSAT is designed to understand performance requirements of structured grid applications, thus can give a tighter bound. In addition, ExaSAT combines software optimizations into the performance model. For more rapid construction of analytic models, pseudo-languages have been proposed. For example Aspen (Spafford and Vetter, 2012) is a domain specific language that enables a user to describe the parallelism, arithmetic operation counts, and data movement to build a model. Specifying the model in Aspen still requires a lot of work, and the quality of the model depends on the ability of the user to accurately capture the application signature.

Our ExaSAT framework has adopted a compiler-based approach to automate the process of generating the performance model. Compiler-based approaches have the dual advantages of being less labor-intensive (thus more easily applied to large codes) and providing more accurate description of codes to the analytic model. We acknowledge that static analysis cannot capture the dynamic behavior of the application; however, metrics gathered by dynamic traces or binary instrumentation are

very sensitive to compiler flags and machine configuration, which can obscure conclusions during the analysis. Moreover, existing machines do not reflect the characteristics of exascale machines and early prototypes of exascale hardware are not available for evaluation. Instead, in ExaSAT we parameterize the hardware configurations to support the static compiler analysis and increase the flexibility of the framework to support exascale machine models. The model is not completely agnostic of inputs either. Rather it is parameterized by a number of runtime parameters such as problem size. These parameters are extracted from the input deck for the various applications so that the model sees the resultant performance impact.

Another aspect that differentiates our approach from others is that our framework uses unconventional performance metrics to quantify performance. [Most existing performance analysis and instrumentation focus on flop counts, cache hit rates, and other processor-based metrics.](#) We focus on data movement as a key metric because it has become one of the most challenging hardware constraints for the design of future systems. Since flops have become cheaper, the energy of data movement dominates the energy cost (Shalf et al., 2010). Thus, our analysis of both on-node and off-node data movement not only provides valuable feedback to hardware designers, but also to exascale programming model, compiler and runtime designers.

Finally, in addition to a parameterized machine model, our modeling approach includes a parameterized model for software optimizations. Previous work (Mohiyuddin et al., 2009; Chan et al., 2013) showed that estimating hardware requirements on an unoptimized software led to incorrect conclusions. Similarly, tuning software without taking into account the hardware choices did not result in an optimal solution. These findings motivated us to incorporate a parameterized set of software optimizations into our framework. Our approach holds a substantial advantage over studies (Balaprakash et al., 2013) that measure code bandwidth and flop utilization without considering software transformations. As more detail emerges on hardware design proposals, the upper bounds provided by the analytic models produced by ExaSAT should be examined together with lower bounds supplied by binary instrumentation on current machines to provide a complete picture of theoretical vs. achievable performance.

Figure 1: Workflow in ExaSAT Framework

3 ExaSAT Framework

As illustrated in Figure 1, the ExaSAT framework is composed of two main components – the front-end compiler analysis and back-end performance model. The front-end component collects procedural and loop level information to create a profile of the code, which is stored in an XML file. The XML code description is then fed into the back-end analysis, which produces dependency graphs, generates performance models, and produces statistical summaries of the code’s characteristics. The performance model is parameterized with (1) machine specifications such as cache size, (2) user parameters such as problem size, and (3) software optimizations such as loop fusion and blocking. The optimized hardware configurations provide the reduced design space for the architecture simulators and the optimized software configurations provide feedback to application developers and programming system tools.

As a result of interactions with application experts and industry partners participating in the DOE Fast Forward program (Fast Forward, 2013), we assembled a set of performance metrics that reflect the characteristics of the exascale applications and made these metrics the center of ExaSAT. Figure 2 shows the list of metrics that we used for evaluating various hardware components and software optimizations. [The first metric that quantifies the benefits of data movement optimizations is the byte-to-flop ratio \(B:F ratio\), which expresses the balance between the application’s required flops and memory traffic.](#) We use this ratio as an indicator of energy and performance improvement for optimization strategies. Since the degree of reuse enabled by the on-chip memory configuration can significantly impact memory traffic, we also measure working set sizes.

Another metric related to data movement concerns state variables and registers. We analyze state

Metric	Corresponding Analysis
Memory Traffic & B:F Ratio	Sensitivity to the memory bandwidth as a result of data movement optimizations
Working Set Size	Data reuse strategies for filtering memory bandwidth
State Variables	Effect of number of registers to avoid register spilling
Arithmetic Operations	FP instruction mix, special hardware, & benefits of vectorization
Read/Write Ratio & Write Access Rate	Candidate streaming data for secondary nonvolatile memory
Fraction of Communication	On-node vs off-node data movement

Figure 2: Subset of performance metrics captured

variables to estimate the impact of the architectural register count on the number of spills, which cause additional loads and stores and pipeline bubbles. Although the majority of our analysis focuses on data movement, ExaSAT provides estimates of arithmetic costs as well. We analyze the instruction mix, the use of expensive transcendental functions such as exponentials, and the impact of vectorization.

ExaSAT enables us to investigate alternative technologies such as non-volatile memory (NVRAM) and integrated network controllers (NIC). NVRAM is considered to be a cost-effective alternative that can serve as a high capacity, secondary memory (Lee et al., 2009), however the writes to NVRAM are costly both in terms of dynamic energy consumption and performance. In order to assess what data to put into NVRAM, we use the community standard, read/write ratios (Li et al., 2012), and a new metric *write access rate*, which is the fraction of write references to a particular variable. Lastly, we use the fraction of communication time to assess the impact of off-node communication on application performance. This metric helps us evaluate whether there is a strong justification for custom NICs, which integrate the network controller on chip to increase injection bandwidth. Next, we explain how we extract these metrics with compiler analysis and how we employ them in performance model.

3.1 Compiler Analysis

The compiler analysis for ExaSAT was built on top of the ROSE compiler framework (Quinlan et al., 2002), which is an open-source compiler infrastructure developed at Lawrence Livermore National

Laboratory. ROSE parses C, C++, and Fortran source to convert it into an abstract syntax tree (AST) that we can manipulate and analyze. Our compiler analysis currently accepts Fortran inputs; however, it can be extended to support C/C++ inputs.

3.1.1 Procedure and Loop Attributes

The analysis of the AST begins by querying the procedure definitions (subroutines and functions) in a module. For each procedure, we collect a list of variable symbols used in the procedure body and classify them into two categories: L : locally declared variable symbols and, U : variables symbols referenced within a procedure. The set difference ($U \setminus L$) of these two gives us the non-local variables, which can be global, defined in another module, or passed as an argument to the procedure. Fortran presents a special case because of its pass-by-reference semantics for subroutine arguments; procedure arguments that are declared with an *intent* type modifier are not technically local, and must be excluded from the L list.

After completing the live variable analysis¹ and locality analysis for the procedure, we collect detailed loop level information. The loop analysis handles perfectly and imperfectly nested loops and is carried out inclusively on the entire loop body without excluding child loops.

When we generate the XML output, we make the loop level information exclusive (not including attributes in the subloops). This is important to accurately estimate performance because both the arithmetic and memory operations are multiplied by the iteration space in the performance model.

For each loop, we gather loop attributes such as iteration bounds and strides, which are later used by the performance model to reason about the iteration space. Loop bounds typically depend on application parameters that are determined at runtime; therefore, we track symbolic rather than actual values and later perform symbolic replacement based on the user’s parameters. Maintaining a symbolic representation of the iteration space also enables the performance model back-end to analyze the effect of software transformations such as loop blocking and fusion.

In order to estimate the total arithmetic workload, we count the floating point arithmetic operations (addition, subtraction, multiplication and division) in the loop body. In addition, we count math intrinsic functions (e.g. *exp()* or *log()*) because they can be significantly more costly to execute. The

¹A variable is live if it holds a value that may be used in the future (thus it cannot be deallocated or overwritten).

compiler analysis searches for function reference expressions in the loop body and uses a lookup table to identify such functions.

3.1.2 Data Access Analysis

Array access analysis is one of the crucial parts of the compiler analysis because the read/write properties of arrays are utilized by the performance model to compute on-chip data movement and memory footprint. ROSE provides an interface to get lists of the read references (R list) and write references (W list) for a given statement. This interface partially serves our needs by enabling us to classify variables as read-only, write-only or both. In scientific codes, some array dimensions may not represent spatial dimensions, but rather different physical properties or quantities such as density, temperature, pressure, or energy. We differentiate such dimensions by representing each array as a array-component pair. For example, the two references to the array Q in $Q(i, j, k, imx)$ and $Q(i, j, k, imy)$ refer to two different array-component pairs: (Q, imx) and (Q, imy) . The location of such dimension is tunable. However, we require all the arrays in the application have the physical property represented in the same index location. The user can identify which indices represent spatial dimensions and which are non-spatial parameters. The differentiation of arrays at the component level is necessary because the reuse pattern, and thus the working set size, is different for each component. We group references by array-component pairs and return separate lists for the read-only variables ($R \setminus W$), write-only variables ($W \setminus R$), and the arrays that are both read and written ($R \cap W$).

In order to model data reuse in the cache, we need more information with respect to the array access patterns. We support the read/write property analysis by examining all the references to an array-component pair in a basic block. The array references are broken into individual subscript expressions to extract their relative offsets to the loop indices. This helps us determine the distance between two references to the same array. Another important property is whether the first reference to an array is a load or a store. If the first reference is a load followed by a store, the load requires the data to be brought into cache from memory before it is written. On the other hand, if a load is preceded by a store, then the load may be carried out from the cache without incurring any additional memory traffic. Our tool conducts the first reference analysis within a loop to more accurately model cache

reuse and support the analysis of advanced memory instructions such as non-temporal stores.

If the program expands into multiple files, we require the user to generate a separate XML per file. The XML code description contains the function call information such as module:function name and argument-parameter mapping. The performance model will take the XML of the file of interest with its dependent XML files as an input. When estimating the performance, if a function call is encountered, the compute cost of that function will be added to the compute time of the callee. Because of the nesting nature of function calls, this complicates the performance analysis particularly for the read/write properties of arrays. We are still investigating how to improve the analysis and currently conservatively assume that arrays are modified if we cannot automatically determine if the function has side effects. We made an exception for the side effect analysis of Fortran intrinsics and assume the arguments for such functions are read-only. In addition to arrays, we conduct a similar analysis for the scalar variables referenced in each loop to help estimate register usage, though the read/write property analysis for scalar variables is much simpler.

3.2 XML Description

The ExaSAT compiler analysis outputs the results in an XML intermediate representation (XML-IR) to interface with the back-end performance modeling component of the framework. This enables the utilization of the performance model directly from the high-level XML-IR, bypassing the compiler analysis step. In this way, program variants or hypothetical code formulations can be evaluated without having to write the actual code. Similarly, the XML output of the compiler analysis can be fed to another tool such as an architecture simulator bypassing the performance model step. We currently provide the XML description of the communication patterns in the codes to the SST simulator (Rodrigues et al., 2011) to simulate different interconnection topologies. A more detailed design document for the XML-IR can be found in (ExaSAT XML Specification, 2013).

3.2.1 Machine Configuration

The machine configuration used as an input to the performance model can be specified in a separate XML. ExaSAT focuses on the aggregate performance of the computational throughput and memory

Figure 3: XML-IR element node hierarchy

bandwidth between the CPU and the DRAM. It considers the bandwidth filtering capability of last level cache, which is determined by the total amount of exclusive on-chips memory per thread or group of cooperating threads. An example machine configuration XML, shown in Listing 1, represents an exascale extrapolation of a many-core architecture. The example shows a 1000 core machine with 10 TF aggregate computational throughput, 1 TB/s aggregate memory bandwidth and 64kB cache per core. The XML also allows us to specify other parameters, such as number of registers, DRAM size, network latency, and network bandwidth. Section 3.3 explains the effects of these properties in greater detail.

```

1  <machine>
2    <prop key="Cores" val="1000" />
3    <prop key="Gflop/s/core" val="10" />
4    <prop key="GB/s/core" val="1" />
5    <prop key="Cache/core (kB)" val="64" />
6    <prop key="Division Cost" val="39" />
7    <prop key="Transcendental Cost" val="125" />
8    <prop key="NIC BW (GB/s)" val="100" />
9    ...
10 </machine>

```

Listing 1: Example XML Machine Description (partial)

Additionally, some software parameters that affect performance, such as the use of cache-bypassed writes and non-temporal memory accesses, may be configured in our performance model through XML input or at runtime.

3.3 Performance Model

Our performance model takes the characteristics of the computational workload specified as an XML and generates performance metrics and execution estimates. For simplicity, we adopt a hardware model abstraction consisting of a collection of parallel hardware cores alongside a parameterized memory on the chip. The CPU is connected to main memory by a bandwidth-limited off-chip network. Our CPU model does not capture the behavior of individual cores or on-chip network, but rather takes the aggregate computational throughput as an input parameter. Similarly, the memory model takes the aggregate DRAM bandwidth connecting the CPU to memory (i.e. the stream bandwidth) as an input. Since modeling the effects of on-chip access latency would require a detailed on-chip network design analysis, we focus on the bandwidth filtering capability of the on-chip memory, i.e. the reduction in memory traffic from capturing temporal locality. Thus, we are primarily interested in the size of the (non-inclusive) on-chip memory capacity per thread or group of threads cooperating on a working set. Our model focuses on capturing the costs of the computational workload and data movement, while taking into account the degree of data reuse enabled by the on-chip memory.

Application performance is estimated using the following method: let α be the aggregate computational throughput of the machine and β be the aggregate memory bandwidth. Let C represent the program’s floating-point arithmetic workload and D be the necessary off-chip data movement between the CPU and DRAM. Our model estimates the program execution time as $T = \max(T_c, T_d)$, where $T_c = \frac{C}{\alpha}$ is the CPU time and $T_d = \frac{D}{\beta}$ is the DRAM time. This performance metric assumes the full throughput and bandwidth are achievable, which may not always be the case for a complex application code. The purpose of our framework is not to make exact performance predictions, but instead provides a performance upper-bound in the spirit of the Roofline model (Williams et al., 2009), and is useful to make relative comparisons between different hardware/software configurations. Lastly, we modeled the off-node communication time by assuming an ideal interconnection network. Our model estimates the communication time as $\frac{m}{b} + l$, where m represents the aggregate message size, b is the network injection bandwidth, and l is the network latency. Thus, for large messages, the network latency is negligible. The model also computes the fraction of communication time over total execution time, which depends on the T_c (on a memory bandwidth limited kernel and T_d (on a compute bound kernel.

3.3.1 Floating-point Computation

In order to estimate C , the floating-point (FP) arithmetic workload, we examine the FP operation distribution present in the code. Current floating-point logic is typically optimized towards FP additions and multiplications, thus exhibit their peak throughput on workloads that only consist of a balance of those two operations. However, there are other types of FP operations present in scientific codes that can only sustain a fraction of the peak. For example, on the Intel Sandy Bridge architecture, the throughput of scalar FP division is 39 times slower than SIMD FP adds or multiplies, while scalar exponentiation is 125 times slower (Vladimirov, 2012). ExaSAT weighs operations such as divides and transcendentals according to their costs specified by the user in the machine configuration to determine a weighted computational workload. Further, the model is parameterized to allow exploring optimizations such as vectorized operations.

3.3.2 State Variables, Registers, and Spills

The number of accesses to both state variables (scalars and non-streamed arrays) and streamed arrays can be used to determine how many registers need to be reserved to hold these values during each of the loops in the program. Since state variables are accessed during every iteration of a loop, an optimal allocation for these variables would place the variables with the most number of accesses into registers, while spilling the rest into the next tier of memory (e.g. L1 cache or local memory). Assuming an architecture with an L1 cache, our performance model can compute the traffic that results from spilled state variables based on the user specified register parameters. In addition, it can compute mandatory traffic that results from streamed variable access to estimate total L1 traffic. This information can be used to analyze the trade-off that results between the number of available registers and L1 bandwidth.

3.3.3 Working Sets and Memory Traffic

The performance model analyzes every array accessed in each loop of the input XML code description. Each array may have a different access pattern, so the tool computes working set and bandwidth usage for each array independently given the array’s access pattern. An array that is written will typically only require access to the current grid element (no neighbors), while arrays that are read

Figure 4: Working sets required for different levels of reuse for a 7-point 3D stencil. The grid is swept in a triply nested loop with the X dimension first, then Y and then Z.

may require multiple grid elements. Our memory and cache model is targeted to the reuse pattern that occurs in stencil computations because stencils constitute the most prevalent operator in our target application codes. The cache is assumed to be an ideal, fully-associative LRU cache, which is optimistic in the sense that if the working set fits into cache, full reuse of that working set is assumed. Real caches with random replacement policies are likely to under-perform due to conflict misses and imperfect replacement. However, our model provides a performance ceiling and a starting point for more detailed analysis using dynamic instrumentation and simulators.

Figure 4 shows the potential reuse cases captured by our model for the canonical 7-point stencil. If the cache is large enough to hold the cell working set, then there will be reuse between cell iterations. Similarly, the figure shows the working set sizes needed for reuse between pencil iterations (all points in x for a given y and z) and plane iterations (all points in x and y for a given z). For each stencil access pattern encountered in the code, our model computes the working set sizes required for each of these reuse cases.

If there are gaps in the stencil access pattern, partial reuse may occur near the calculated boundaries between reuse cases. Our model can compute the working sets sizes that bound the transitions from no reuse to partial reuse to full reuse between pencil and plane sweeps. For an LRU cache, no reuse will occur if the cache is smaller than the number of elements accessed in the pattern, while full reuse requires a working set equal to the span of the pattern plus the maximum gap size. For example, a stencil pattern that accesses planes $-2, -1, 0, +1, +2$, has a working set of 5 planes because there's

no gap, but a pattern of $-2, +2$ requires a working set of 8 planes (5 for span, 3 for gap) for reuse even though it only touches 2 planes per sweep. It may seem counter-intuitive that accessing fewer planes can increase the working set size, but gaps in the pattern require the cache to hold data for a longer period of time without evicting it. For a software-managed local store, the memory can be managed more efficiently, requiring only the span of the access pattern to fit into the store. Since we are interested in establishing a performance upper bound, the model optimistically assumes full reuse is possible for certain situations where only partial reuse would occur. Future work will take these effects into consideration to increase the tightness of the bound.

The machine configuration specifies the cache line size, which determines the minimum granularity of access in the unit-stride dimension used for working set and bandwidth calculations. Our model rounds the number of contiguous elements within the accessed region up to the next multiple of the cache line size (assuming optimistic alignment), to compute the resulting working set and memory traffic estimates. Also, the configuration allows the user to specify whether cache bypass is utilized for array writes, reducing memory traffic and cache pollution. Non-temporal array reads can also be enabled in the configuration to further reduce cache pollution from arrays with no reuse.

Once the working set and memory traffic estimates are computed, they are compared to the cache size specified in the hardware configuration to determine what reuse scenario will occur for each loop, thus determining the required memory traffic for the whole program. Note that this type of analysis can be conducted at every level of the cache hierarchy. For example, if we specified the cache size available at L1, then the computed memory traffic would be that required between L1 and L2. Using our methodology, we could conduct a multi-level analysis that computes the bandwidth requirements and performance at every level of cache.

3.3.4 Block Execution Schemes

Cache blocking (Rivera and Tseng, 2000) reduces cache capacity misses by tiling the loop iteration space, thus shrinking the working set to the point where it fits in cache. ExaSAT incorporates two different block execution schemes to analyze the performance impacts of cache blocking. In the traditional blocking scheme, each loop runs over the entire domain before proceeding to the next loop.

Figure 5: Comparison between a) traditional blocked execution order and b) the alternative block execution order

In an alternative scheme (Woodward et al., 2010) all of the loops are run on a block before moving to the next block as illustrated in Figure 5. Each large rectangle represents the iteration space at different points of progress (indicated by shading), and each sub-rectangle represents a block of the iteration space that fits into local memory. While traditional blocking allows reuse of data within loop nests, the alternative scheme schedules loops such that reuse of data *across* loop nests is also possible. The potential disadvantages of the alternative scheme are larger working set sizes and redundant overlapping computation needed to satisfy any necessary spatial dependencies between blocks. If the blocks are sized appropriately, all temporary arrays can remain in cache or local store throughout the computation until the final output is produced. If there is sufficient on-chip memory, the only DRAM traffic required would be for reading and writing each function’s inputs and outputs.

ExaSAT automatically generates parameterized performance models for both schemes, facilitating the exploration of optimal strategies for different machine configurations in the co-design process. Using liveness analysis, our performance model can estimate the total memory footprint needed at each computation step, giving the on-chip memory size required for each block execution scheme and an estimate of the total memory traffic.

3.4 ExaSAT Outputs

3.4.1 Dependency Graph Description

The ExaSAT framework outputs a dependency graph indicating the dependencies between loops in a procedure. Flow, anti, and output dependencies are considered across all arrays read and written in

Figure 6: Dependency subgraph for the SMC dynamics code

Table 1: Example loop analysis table (top) and array access and occupancy table (bottom) generated by the ExaSAT tool for a subset of SMC dynamics code

Procedure	Loop Line Number	Add	Flop's/cell Mul	Div	Exp	State Int	Var FP	Working Set (kB)	Memory Traffic (GB)	FP Computation (Weighted GFlop's)	B:F Ratio	Execution Times (ms)
advance	418	128	174	0	0	17	5	356	0.69	0.63	1.16	0.67
advance	533	2	2	0	0	7	2	0.03125	0.11	0.01	12.00	0.11
advance	720	32	39	0	0	13	8	132	0.19	0.15	1.39	0.19
advance	771	18	27	9	0	17	0	0.4375	1.41	1.00	1.52	1.37
advance	1529	860	959	18	0	30	70	818	1.44	5.33	0.29	1.41
ctoprim	85	3	17	1	0	24	32	0.375	1.54	0.15	11.12	1.50
ctoprim	136	4	4	2	1	14	22	0.1171875	0.23	0.44	0.57	0.23
Total/Max								818	5.61	7.71	0.78	5.48

Variable Name	Copies	Loop Line Number												Totals		
		274	418	515	767	771	791	1139	1160	1508	1529	1877	1921	Reads	Writes	Live
Fdif.iry	53				W	L	L	RW	L	RW	L	RW	R	212	212	265
Fhyp.iry	53	W	RW		L	L	L	L	L	L	L	L	R	106	106	477
Hg.iry	53						W	R	W	R	W	R		159	159	0
Q.qhn	53					R	R	L	R	L	R	L	R	212	0	106
Q.qpres	1		R	L	L	R	R	L	R	L	R			5	0	4
Q.qtemp	1						R	L	R	L	R			3	0	2
Q.qxn	53					R	R	L	R	L	R			212	0	106
U.iry	53	L	R	L	L	L	L	L	L	L	L	L	R	106	0	530
Unew.iry	53	L	L	L	L	L	L	L	L	L	L	L	RW	53	53	583
dpe	1					W	RW	R	L	R	L	R		4	2	2
dpyn	53						WR	R	L	R	L	R		159	53	106
Number of Arrays Resident		159	160	213	214	373	427	427	427	427	427	265	212			

each loop. Figure 6 shows an example dependency graph generated by ExaSAT where boxes represent data arrays, ovals represent loops, and arrows indicate which arrays are read and written by each loop. The dependency graph illustrates the code’s inherent concurrency and allows us to reason about how the computation can be rearranged for enhanced locality, task co-scheduling, and parallel load distribution. We explore the impact of loop fusion for enhanced locality on our motivating application in Section 5.1.4. Future work will study the use of intelligent runtime analysis for task co-scheduling and load balancing.

3.4.2 Spreadsheet Description

ExaSAT outputs a performance spreadsheet for the user to further examine the performance of the application. The spreadsheet contains a table of user-modifiable parameters, which allows the user to change the initial XML software and machine configurations. The rest of the spreadsheet automatically updates itself via formulas to reflect the changes made in the parameter table.

The main section of the spreadsheet is a summary table listing properties for each loop in each procedure in the code. Table 1 shows a part of the summary table generated by our tool, including flop counts, state variable count, working set size, memory traffic, and execution time. Aggregate statistics are also included in the table to summarize whole program performance.

The spreadsheet contains an array access and occupancy table which shows the liveness of arrays through the progression of the program. This analysis is used for memory capacity calculations and NVRAM feasibility studies. Table 1 also shows an example occupancy table generated by our tool. The rows in the table correspond to arrays, while the columns correspond to loops in the program, allowing the table to be read left-to-right to correspond to a possible program execution. The number of copies indicate the number of components of the arrays. Each cell in the table contains one of the following values: read (R), written (W), read-then-written (RW), written-then-read (WR), live (L), or non-resident (). Summary columns are given to show the number of reads and writes to each array as well as the total number of live arrays, which is used to compute the memory footprint.

Other sections in the spreadsheet include tables that summarize the inter-node communications that must occur during the program execution and state variable accesses to help model the cache traffic

Figure 7: Comparing memory traffic modeled by ExaSAT and simulated by Pin for the CNS code (128^3 problem)

resulting from spilled registers. A summary table and histogram are generated for each loop showing the number of state and streaming variables located in registers versus cache and the corresponding number of register hits and misses.

3.5 Model Validation

We validated our results against data collected through dynamic instrumentation and benchmarking.

Validation against Pin: First, we used the publicly available Pin tool (Luk et al., 2005) to validate instruction counts and memory traffic. Pin analyzes an application at the instruction level and uses dynamic compilation to instrument executables while they are running. By attaching callbacks around every instruction reading or writing to memory we can extract a stream of load and store addresses from the program as it runs. This stream is then piped into a LRU cache simulator that we implemented on top of Pin, which aggregates the relevant statistics such as cache hit, miss, and line writebacks for a given cache size. Floating point instructions are also monitored to retrieve flop counts.

Figure 8: Measured and modeled execution times for blocking sizes for three simple stencil benchmarks

The FP instruction counts predicted by ExaSAT match with those measured by the Pin tool. Loads and stores of the variable under study also match those reported by the Pin tool. Figure 7 compares the memory traffic modeled by ExaSAT with the memory traffic captured by the Pin tool for the CNS code for various cache sizes. CNS² is a combustion proxy (Exact, 2013) that integrates the compressible Navier-Stokes equations assuming constant transport. It is a simplified (single species) version of the SMC code (Emmett et al., 2013), which will be discussed in more detailed in Section 4. The analytical performance model in ExaSAT correctly captures the amount of data reuse and resulting trend of memory traffic as cache size is varied, though the memory traffic modeled by ExaSAT is slightly lower than the Pin tool’s because it is providing a lower bound. Initially, the number of L1 cache hits measured by Pin was abnormally higher than what ExaSAT estimates considering only array access traffic, which led us to investigate the proportion of L1 cache traffic due to spilled state variables. When there is not enough number of registers to hold all the state variables in a loop, accesses to these variables will be spilled to the next level memory. This introduces more cache traffic, which will give the impression that there is a higher hit rate. When we separated array accesses from the state variable accesses to the cache in the Pin tool, then the loads and stores estimated by ExaSAT match with those measured by Pin. More discussion on register spills will be described in Section 5.1.3.

Block Size Validation: Second, we measure the effect of blocking with three simple stencil benchmarks, namely gradient, divergence and Laplacian and compare their performance against the

²CNS is available for download at the ExaCT co-design center’s website (Exact, 2013).

Figure 9: Measured and modeled execution times for each loop in a single Runge-Kutta step in the CNS code. $L\#$ indicates the loop numbers.

estimates by ExaSAT. We manually blocked three simple stencil benchmarks, namely gradient, divergence and Laplacian and collected the execution time with 24 threads on a single node on NERSC Hopper Cray XE6. No software prefetcher or cache bypass is enabled. The results in Figure 8 show that the measured execution times and optimal blocking size correlate well with ExaSAT's. Where the block size is small, the model predicts much better performance than the measured because in the measured code, hardware prefetchers cannot hide the load latencies for small blocks. There are also situations that the model exceeds the measured execution time. The model has a sharp transitions at the points where the working set grows larger than the available cache. In reality, the cache replacement policy leads to a smoother transition than ExaSAT.

Optimization Opportunities: Third, we collected running times for the CNS code (single species) and SMC code (multiple species) and compared them against the ExaSAT estimated bounds. The purpose of comparing ExaSAT with the benchmark is not to measure how close the estimated

and actual running times are but to point to the parts of the code where there are opportunities for optimization since ExaSAT highlights parameter sensitivities subject to the user-specified constraints rather than giving a performance prediction.

Figure 9 compares the performance bounds by ExaSAT and the actual running times by loop nests collected on the NERSC Hopper machine³. As clearly seen from the results, some of the loops have a big performance gap between the two and these have the potential to gain some of the performance back through data movement optimizations. In particular, the *hypterm* function exhibits the highest discrepancies between the measured and estimated bound. ExaSAT bounds the running time for the *hypterm* function to 8.1ms, which is 3.3x better than what is measured (26.6ms). ExaSAT estimates that this bound for *hypterm* can be further reduced to 3.3ms from 8.1ms if all the three loops are fused.

We aggressively optimized the *hypterm* function by applying vectorization, cache blocking, and loop fusion optimizations and the results are shown in the inset graph in the same figure. The first bar in the inset graph shows the total time spent in three loops in *hypterm* and the second bar shows the measured performance as a result of optimizations and compares it with the new bound by ExaSAT after the loop fusion. The manual optimizations achieved 6.7x the initial performance, reducing the measured running time from 26.6ms to 4.0ms, which is much closer to what ExaSAT predicts (3.3ms). This illustrates how ExaSAT can be used to identify performance opportunities for the programmer and guide application tuning.

Similarly, ExaSAT suggested that tiling the SMC code would provide 37% improvement in performance on Hopper and 41% speedup on SDSC’s Trestles⁴. We have implemented a tiled version of SMC and observed a 30% improvement on Hopper and 32% on Trestles. We suspect that the lower measured performance can be attributed limitations with the hardware prefetchers since the SMC code accesses a large number of arrays in its solvers. Consequently, there is room for improvement and we are still investigating the SMC performance. We did not manually implement the fused version of SMC because of its complexity. We plan to use the Chill compiler framework to automate loop fusion.

³using six core Opteron 6172 with 6MB L3 cache.

⁴using four AMD MangyCours with 4MB L3 cache.

4 Motivating Application: SMC

We demonstrate the abilities of the ExaSAT framework by applying the tool to the SMC code, which contains over 10K lines of code, making the manual analysis impractical for this code. SMC is developed by the Combustion Co-design Center (Exact, 2013) and is a proxy for the production direct numerical combustion codes such as S3D (Chen et al., 2009). SMC represents structured grid problems, which play an important role in numerical simulations, particularly in stencil-based PDE solvers. Understanding the SMC performance provides insights into requirements of the family of combustion codes on exascale machines.

SMC integrates the multicomponent reacting compressible Navier-Stokes equations with detailed models for chemical species diffusion and kinetics. It contains the key elements of both the dynamical core⁵ and the chemical kinetics components of S3D; however, SMC is restricted to gas phase problems and a restricted set of boundary conditions. SMC also uses a simpler temporal integration algorithm that does not include automatic error control. The methodology is based on high-accuracy solution of a system of partial differential equations of the form

$$\frac{\partial U}{\partial t} + \nabla \cdot \mathcal{F}(U) = \nabla \cdot \mathcal{D}(U) + \mathcal{S} \ .$$

The terms $\mathcal{F}$, $\mathcal{D}$, and $\mathcal{S}$ correspond to hyperbolic transport, nonlinear diffusive processes and chemical source terms, respectively. U is a vector of unknowns, representing density, energy and three components of momentum with an additional density for each chemical species (e.g. octane), for a total of $5 + N_s$ unknowns per point where N_s is the number of species in the problem. The number of chemical species and the number of reactions have a strong effect on overall computational costs of the algorithm; typical applications will range from as few as 9 species to more than 100. The chemical kinetics model used by SMC is specified at compile time using code that is generated automatically from a tabular description of the reaction mechanism. This mechanism-specific file also includes thermodynamic data needed for the simulation. Transport coefficients are computed using EGLIB (Ern and Giovangigli, 1995).

⁵The part of the code that computes fluid dynamics.

We focus on two important aspects of SMC: the chemical source term evaluation and the dynamical core. The chemical source term of SMC is a computationally intensive, element-wise computation that uses a large number of transcendental operations. The dynamical core uses high-order stencil computations to approximate spatial derivatives, converting the system into a large system of ordinary differential equations. These ordinary differential equations are then integrated using a third-order, low-storage, TVD Runge-Kutta scheme (Gottlieb and Shu, 1998; Qiu and Shu, 2005).

The spatial discretization uses a finite difference approximation on a uniform grid. There are essentially three types of terms we need to approximate: first-order derivatives needed to approximate $\nabla \cdot \mathcal{F}$ and terms of the form $(au_x)_y$ and $(au_x)_x$, both of which arise in discretizing $\mathcal{D}$. We first define a first-order derivative operator in the x direction, $D^{1,x}$ using an eighth-order finite difference discretization

$$u_{x,i,j,k} \approx D^{1,x}u_{i,j,k} = \sum_{\ell=1,4} \alpha_{\ell}(u_{i+\ell,j,k} - u_{i-\ell,j,k})$$

with analogous operators in the y and z directions. These discrete derivative operators are used to evaluate the terms for discretization of $\mathcal{F}$. They are also used to evaluate mixed derivative terms. For example,

$$(\eta u_x)_y \approx D^{1,y}(\eta D^{1,x}u) .$$

The second derivative terms are discretized using an eighth-order extension of the narrow stencil discretization of Kamakoti and Pantano (Kamakoti and Pantano, 2009). In particular, we approximate variable coefficient second derivative terms in the form

$$\frac{\partial}{\partial x} \left(a \frac{\partial u}{\partial x} \right)_i \approx D^{2,x}(a, u) = \sum_{\ell,m=-4,\dots,4} \beta_{\ell,m} a_{i+\ell,j,k} u_{i+m,j,k}$$

A more detailed discussion of the discretizations in SMC can be found in (Emmett et al., 2013).

The parallel grid decomposition for the SMC code requires *ghost cell* exchanges for the vector U . Ghost cells are the data residing in neighboring grid blocks that are required to compute the stencil operations. Figure 10 shows the stencil access pattern and ghost region that needs to be communicated. The depth of the ghost region is four grid cells in each dimension with $5 + N_s$ values per point.

Figure 10: (a) 3D grid with its ghost cells (b) stencil access pattern for SMC

	Description	# of Species	# of Reactions	# of Reactions/Species
LiDryer	Hydrogen	9	21	2.3
Drm19	Reduced reaction sets natural gas	21	84	4.0
Grimech30	Natural gas combustion	53	325	6.1
Hai	Tri-carbon fuel combustion	71	469	6.6
Prf_ethanol	Ethanol	107	529	4.9

Figure 11: Number of species, number of reactions and the description of chemistry component of the SMC codes modeled

5 Results

5.1 Analysis of SMC with ExaSAT

In order to capture the effect of increasing number of species, we modeled the SMC code for 9, 21, 53, 71 and 107 species, representing simulations ranging from hydrogen to natural gas to biofuels. Figure 11 provides more descriptions about the species modeled in this paper. Note that only certain values for the number of species are meaningful. Unless stated otherwise, the baseline performance estimates are based on a domain decomposition (box) size of 128^3 per node with 53 species using the machine configurations specified in Listing 1.

5.1.1 Arithmetic Operations

Scientific applications are biased heavily towards floating point operations, although address arithmetic is often a large component of the instruction mix. The preponderance of floating point operations are

Table 2: Relative throughput of divide and exponential compared to vectorized ADD on Intel Sandy Bridge E5-2680 with Turbo Boost

Relative throughput	Baseline	Fast-div	Fast-exp
Division	1/39	1/20	-
Exponentials	1/125	-	1/42

addition and multiplication operations. A small number of division and transcendental functions that appear in the codes contribute significantly to the running time since they execute one or two orders of magnitude slower.

ExaSAT can examine the floating point(FP) operation mix per loop iteration and provides the flexibility to change the arithmetic operation throughput. Figure 12 shows the operation analysis for both the chemistry and dynamics kernels of the SMC code for a 128^3 problem size with 53 species. The two kernels exhibit substantially different arithmetic operation distributions. The chemistry kernel contains transcendental operations, mainly exponentials (92.5%) and logarithms (7%). Even though, a small number of division and transcendental functions appear in both kernels, these operations contribute significantly to the running time since they execute roughly one to two orders of magnitude slower. Figure 12 shows the estimated contribution of each FP operation to the CPU time when we assume vectorized addition and multiplication, and low throughputs⁶ for division and transcendental functions (1/39th and 1/125th of peak, respectively). Note that the CPU time (T_c) is computed based on the compute throughput and does not include the DRAM time. Even though transcendental functions in the chemistry kernel are a small fraction of the total flops, they dominate the CPU time (75%). Similarly, the number of divisions in the dynamics kernel seems insignificant but contributes one third of the CPU time.

5.1.2 Fast Transcendentals and Division

Vectorization is one of the main sources of parallelism within a processor that can enable fast execution of floating-point division and transcendental arithmetic. Besides parallelism benefits, it can also lower energy and control complexity. The downside is that it takes chip surface area and requires programmer assistance. An alternative approach to vectorization is software pipelining, which can hide functional

⁶Based on the benchmarks we conducted and by (Vladimirov, 2012).

Figure 12: Floating point operation mix and breakdown of CPU time T_c modeled by ExaSAT for chemistry and dynamics kernels.

unit latency but also requires more programming effort and more registers.

Table 2, shows benchmarked (not modeled) performance results gathered on the Intel Sandy Bridge E5-2680. *Fast-div* shows the performance improvements for division using the SSE instruction (AVX provides no further performance gain) and *Fast-exp* shows the performance improvement for the exponential function with the AVX Short Vector Math Library. The benchmark results indicate that SSE provides 1.95x improvement on division and AVX provides 2.98x improvement on exponentials.

ExaSAT allows a user to weight instructions based on their relative throughput to the peak compute rate. The weights can reflect the longer execution times of certain instructions such as division, or they can reflect the potential speedup derived from more pervasive use of intrinsics through improvements to the compilers or hardware. The speedup due to the use of intrinsics may not be proportional to the increased weight of the instruction speed because the compiler might fail to generate code that uses intrinsics due to the complex loop body or divergence effects.

Figure 13 shows the estimated speedup for the SMC code including both the chemistry and dy-

Figure 13: Estimated overall SMC speedup over baseline as a result of simdizing division and transcendental functions using different vector lengths. Baseline indicates vectorized addition and multiplication operations running with at the peak compute throughput but no vectorization for division or transcendentals.

namics codes as a result of different SIMD lengths. Here, the baseline performance assumes SIMDized addition and multiplication, and low throughputs for division and transcendental functions (first column in Table 2). Figure 13 also shows the estimated speedup for SMC on the Sandy Bridge (indicated as a line) using the benchmarked costs for division and transcendentals shown in Table 2. Our performance model takes the maximum of CPU time and DRAM time for each kernel independently to compute the execution time. Both vectorized division and transcendentals greatly improve the execution time of the chemistry code; however, there is no benefit for the dynamics code since its execution time is limited by DRAM bandwidth. As a result, there is a diminishing return as we increase the SIMD length. For example, for 53 species, the SSE instruction (SIMD2) provides 25%, while SIMD4, SIMD8 and SIMD16 give 43%, 54%, and 60% improvement over the baseline, respectively. The improvement differs between different number of species because of the number of reactions, thus the number of divisions and transcendentals differ. Both 53 and 71 species have a high number of reactions per species, which means more arithmetic operations and higher benefit from vectorization for the chemistry component.

Figure 14: Number of accesses for each floating point state variable sorted by their access frequency in the chemistry kernel

5.1.3 State Variables

The state variable analysis provided by ExaSAT is valuable in the co-design process because it exposes a hardware trade-off between register count and L1 cache traffic (or local memory traffic). In order to measure how many registers the SMC code requires, we collected all the state variables and their access frequencies for each loop using compiler analysis. Based on the number of registers specified by the user, the performance model allocates the state variables to available registers and computes the L1 cache traffic resulting from the register spilling. Figure 14 shows the number of accesses for each floating point state variable sorted by number of accesses in the SMC chemistry kernel. For example, in a 9-species simulation, the variable # 22 is accessed 15 times. In the best case scenario, the compiler will allocate the variables with the highest number of accesses to the available registers. Assuming 16 floating point named registers (as in SSE or AVX), the vertical dashed line shows the cut-off between variables that would be allocated to registers (left of the line) and those that are spilled to cache (right of the line).

Figure 15 shows the percent of state variable accesses spilled to the next level memory as the number

Figure 15: L1 cache traffic chemistry state variables as number of the registers is varied. Having more registers can filter cache traffic for state variables.

of available registers is varied. In the 16 register example, about half of the accesses are fulfilled from registers and half go to cache for each of the five chemistry species shown. In the dynamics kernel (not shown in the figure), even though the total number of state variables is much smaller, assigning the top 16 variables to registers only reduces the number of cache accesses by about half since access rates remain fairly high for the top 30 to 40 variables for many loops. Since the chemistry code has a relatively low streaming data requirement compared to the dynamics code, spilled state variables make up greater than 95% of the L1 cache traffic if there are 16 registers. It is possible to filter additional cache traffic by adding registers to the architecture, which would move the cut-off line in Figure 14 to the right. Having 256 registers per thread (as in NVIDIA's Kepler GPU⁷) would filter 88% or more of L1 cache traffic due the state variable for the SMC chemistry code, and 94% or more for the SMC dynamics code. It is important to note that the spills must be balanced against performance cost of large register file. The optimal performance point may be reached at an earlier point.

⁷Kepler has 255 32-bit registers.

Figure 16: Optimal blocking factor depends on the available cache size.

5.1.4 Memory Traffic and Working Sets

Cache Blocking: ExaSAT can model the effect of cache blocking on the working set size and memory traffic without manually implementing this optimization. Blocking the iteration space shrinks the size of the working set to enable temporal data reuse. If the reduced working set fits within the available on-chip storage, capacity misses, thus the memory traffic can be greatly reduced. A trade-off of cache blocking is the induced memory traffic for the ghost cells. As the block size is decreased, the redundant traffic to pull the ghost zone increases. Finding the optimal blocking factor on a given cache size is an optimization problem for compilers, auto-tuners and runtime environments. In this context, ExaSAT can guide other programming tools to reduce to search space for blocking factor. We are also interested in illustrating the co-design trade-off of blocking, more specifically the trade-off between cache size and memory bandwidth. For a given cache configuration, ExaSAT can determine a blocking strategy that balances the capacity misses against the additional traffic for the ghost zone.

Figure 16 highlights the change in byte:flop ratio of the dynamics kernel computed by ExaSAT as a result of blocking for various cache sizes specified by the user. [The byte:flop ratio represents the required number of bytes to be transferred off-chip divided by the required flops.](#) The cache size

Figure 17: Various software optimizations modeled by ExaSAT for the 53 species SMC dynamics code and their byte-to-flop ratio achievable for various on-chip memory sizes

indicates the amount of on-chip memory available per group of threads/cores collaborating on the same working set. Blocking the iteration space reduces the working set size and enables greater reuse. The inflection points in the plot show the points when the working sets no longer fit into the cache.

However, the trade-off for using smaller block sizes is additional memory traffic from redundant ghost cell storage and accesses. This effect can be seen even in the unlimited cache case because it is independent of capacity misses. Thus, ExaSAT predicts that blocking with an inappropriate factor could incur more data traffic than necessary. With an optimal blocking factor, a small cache can beat the performance of an unblocked reference implementation on a large cache. Consequently, the compiler or auto-tuner has to find the optimal block size to take full advantage of available cache, while chip designer has to find a balance between the cache size and memory bandwidth.

Software Optimizations: ExaSAT also allowed us to evaluate the performance impact of software optimizations such as loop fusion and the alternative block execution scheme described in Section 3.3.4. Even though loop fusion can reduce memory traffic, it increases the resulting loop's working set, exposing a co-design trade-off between memory bandwidth and cache size. Loop fusion was done by

hand, guided by the data dependency graphs generated by the framework, while the cache blocking and alternative block execution schemes were computed automatically from the XML code description. Figure 17 shows effects of applying various software optimizations on the trade-off space between cache size and the resulting B:F ratio. For small cache sizes, no blocking is used, but there is still some benefit from applying loop fusion to loops that touch the same data. For medium cache sizes, some loops are able to take advantage of reuse within loops in the non-fused case, but there is not enough cache to hold the increased working sets required for fused loop bodies. Once the caches are large enough to contain the increased working sets of the fused loops, fusion becomes beneficial again. For 53 species, the breakpoint is about 2MB.

The alternative block execution scheme requires the largest working sets because an entire block of data per array must fit in cache (as opposed to a small number of planes per array) to enable reuse *across* loops. However, the benefit from such reuse is a significantly lower byte-to-flop ratio (roughly half for the largest cache sizes in the figure). This execution scheme may be most relevant to situations with processing capabilities co-located with large memory banks such as with Processor-in-Memory and Processor-Near-Memory architectures (Saulsbury et al., 1996). The studied optimizations emphasize the power of software transformations on the byte-to-flop ratio and their relation to cache size. Not surprisingly as we increase the number of chemical species, the working set size increases (not shown in the figure), requiring a larger cache for fusion to become advantageous. Please see (Chan et al., 2013) for a further analysis of software optimizations on combustion co-design.

5.1.5 Memory Footprint Analysis

ExaSAT can compute the memory required for an application. For SMC, the memory requirement increases linearly with the number of species. A 53 species simulation needs 678 three dimensional arrays, translating into approximately 13 elements per species per grid point. Out of 678 arrays, 505 of them have a ghost cell region. Including message buffers, a box size of 128^3 occupies 12.33GB of memory, which means a 16GB node can only hold one 128^3 box. In order to minimize the surface to volume ratio of each node’s subdomain, there is a motivation to run as large a problem as possible on a node. Hence, combustion codes that use adaptive mesh refinements and implicit time stepping schemes

Figure 18: Read/write ratio (left axis) and write access rate (right axis) of arrays in SMC code

exhibit huge workload diversity in chemical reactions rates, requiring multiple subdomains running on a node for load balance. Exascale memory capacity is predicted to be primarily constrained by cost (Kogge et al., 2008) which encourages vendors to look for cheaper but denser memory technologies such NVRAM. NVRAM is a cost effective alternative technology that can serve as a high capacity, secondary memory. It offers higher density and scalability than DRAM, and uses nearly zero power when in standby mode (Lee et al., 2009). On the other hand, the NVRAM memory cells tend to have a short lifetime. Compared to DRAM, the dynamic write energy is 4 to 40 times worse and the write access latency is an order of magnitude slower (Caulfield et al., 2010; Lee et al., 2009; Qureshi et al., 2009). Nevertheless, without focusing on the details of the NVRAM designs, we investigate whether there is sufficient low-write memory traffic for certain variables to justify inclusion of NVRAM in an exascale node since the specifics of NVRAM designs regarding memory endurance, write-voltage and write speed are highly dependent on the technology and are likely to change.

In order to study the NVRAM opportunities in the application, ExaSAT computes the read/write ratio and write access rate of arrays since writes to NVRAM are costly both in terms of performance and energy. In SMC (see Figure 18), there are a number of arrays with low read and low write access rates. We are primarily interested in arrays rather than scalar variables because the idle power

Figure 19: Fraction of communication time for different Memory Bandwidth/NIC Bandwidth ratios. $\delta = 10$ is an expected value for an exascale node. $\delta = 2.5$ represents a relatively fast network bandwidth and $\delta = 40$ represents a relatively fast memory.

consumption is proportional to the memory footprint. If a write access rate of $\leq 0.11\%$ is chosen, then a larger fraction of data (75%) qualifies for storage in NVRAM. This would translate into roughly 75% idle power saving. On the other hand, the the dynamic energy for these arrays would go up by a factor of 40. Even if a conservative read/write ratio of 5 or higher were chosen, the case for NVRAM would be weak because only 35% of the data would reside in NVRAM. Unfortunately, this is where our analytic model has its limits. To assess whether the dynamic energy consumption overshadows the idle energy savings, power simulators such as NANDFlashSim (Jung et al., 2012) are needed, which is a part of our future work.

5.1.6 Communication Analysis

The interprocess communication time as a percentage of total execution time depends on the DRAM time on a memory bandwidth limited kernel (or CPU time on a compute bound kernel). Figure 19 shows the fraction of communication time for the bandwidth limited SMC code as the memory bandwidth to network bandwidth ratio, δ , is varied. For example, a configuration with 1 TB/s of memory bandwidth and 100 GB/s of NIC bandwidth would correspond to $\delta = 10$, which is an expected value at exascale. The figure varies δ from 2.5 (a relatively fast network bandwidth) to 40 (a relatively fast memory). According to the analytic results shown in Figure 19, communication time accounts for less than 13% of the total time in the worst case and diminishes as we increase the number of species for

the SMC code. Because the communication time does not appear to be a performance bottleneck for SMC, there is no strong justification for integrating NIC on the processor chip to increase the injection bandwidth and reduce latency. Future work will study adaptive mesh refinement codes, where messages tend to be smaller but more frequent. In those cases, we expect there might be more need for on-chip NICs.

The analytic performance analysis is agnostic about network topology and assumes an idealized network for off-node communication, considering only network latency and injection bandwidth as the performance metrics. However, factors that are hard to capture in an analytic model such as network topology, routing, network contention, and job placement can have a significant impact on performance. We are currently collaborating with Sandia National Laboratory to employ the SST/macro simulator (Rodrigues et al., 2011) to assess the network performance of SMC. ExaSAT serves as a stepping stone for such effort and is used to verify the simulation results. For example, the 3D torus with a optimal job placement matches with the idealized network scenario by our analytic model, seeing no network congestion with pure nearest-neighbor communication.

6 Discussions

6.1 Projections on an Exascale Machine

Figure 20 shows the cumulative effect of the hardware and software improvements modeled by ExaSAT. The estimated effective baseline performance is slightly over 0.5 Tflops using the machine configurations specified in Listing 1, which is a 10 Tflop node with 1 TB/s memory bandwidth. The SMC code is severely limited by memory bandwidth. Both cache blocking and loop fusion make more efficient use of memory bandwidth, doubling the baseline performance. However, the estimated performance indicates that software optimizations must be supported by hardware improvements at the expense of increased cost and power for the sake of higher performance. If the memory bandwidth is increased from 1 to 4 TB/s, ExaSAT suggests a 2.5-3x speedup in the performance is possible. We also modeled the effect of vectorization of division and exponentials for the SMC code. *Fast-div* represents the predicted performance improvements as a result of improved throughput (2x) using the SSE instruction

Figure 20: Modeled SMC performance as a result of successive hardware and software optimizations and *Fast-exp* represents the improvement for the exponential function by a factor of 3 with the AVX SVMML. While the vectorized division provides a modest performance increase, the chemistry component greatly benefits from the improved cost for exponential. Finally, we change the network injection bandwidth from 100 to 400 GB/s, which represents a custom NIC. Custom NIC integrates the network controller onto the chip to reduce power and increase throughput by a factor of 4. Even after the software optimizations and hardware improvements, SMC is still limited by memory bandwidth. In the exascale timeframe, it is unlikely that machines will support bandwidths higher than 4 TB/s, thus more aggressive software optimizations will be needed to reduce data movement and deliver the performance improvements necessary to reach exascale.

6.2 Implications for Hardware Design

We evaluated the impact of utilizing vector intrinsics for division and transcendental functions and realized that they can greatly improve the CPU-bound chemistry code provided that compilers or code generators can support vectorization. On the other hand, SIMD lengths more than four do not provide significant performance benefits because the dynamics part of the SMC code is severely limited by the

performance of the memory subsystem. For the baseline code with optimal cache blocking, we see very little benefit derived from larger on-chip caches. However, if we adopt a more aggressive approach with loop fusion, we can achieve an order-of-magnitude reduction in memory bandwidth requirements provided there are much larger on-chip memory and register files. For SMC, fusion can reduce traffic by up to 60% versus baseline provided that there is large enough cache. Having 256 registers per thread would filter 88% or more of the register spills due to the state variables.

In our assessment of data accesses, given that technology allows NVRAM write performance to improve, we see some opportunities to utilize NVRAM to increase memory capacity with low cost. However, the NVRAM technology has to mature before an investment in software support can be justified. In order to determine which data to place to NVRAM, we argue that write access rate rather than read/write ratio should be used as a metric. There is a modest performance benefit from the reduced latency and increased bandwidth of on-chip NICs because the network injection bandwidth does not appear to be a performance bottleneck, and the SMC application is insensitive to interconnect latency.

6.3 Implications for Software Design

ExaSAT is a lightweight model, which can be integrated into other tools and serve as a cost model. In particular, our analysis of data movement both on-chip and off-chip provides valuable feedback to application, programming model, compiler, and runtime developers. The results emphasize the importance of reduction in memory traffic both for performance and energy reasons. In fact, one of the co-authors of this paper implemented a new blocking optimization in SMC based on our ExaSAT analysis, which yielded an 86x speedup over 1 thread on a 61-core Xeon Phi each running 4 hardware threads (Emmett et al. (2013)).

ExaSAT can also provide performance ceilings for compute-bound kernels. Simply having vector units on the chip is not sufficient to increase performance because the compiler also has to generate the appropriate instructions. Current compilers can convert scalar codes into SIMDized programs with some programmer assistance, such as ensuring address alignment and providing compiler directives. Automatic vectorization often fails for complicated loops because other code optimizations may inter-

interfere with vectorization or the loop body may be too long to analyze. The highly irregular structure and single-point implementation of the chemistry code currently prevent the compiler from inserting vector intrinsics, especially on GPU-like architectures. The ExaSAT results encouraged the SMC developers to restructure the chemistry component in way to facilitate vectorization by the vendor compiler and resulted in 2.2x faster chemistry on the Edison machine that include 256-bit SIMD (AVX) vector floating point. The chemical reactions in the SMC code were previously auto-generated in order of species appeared in the input file. Two reactions which have the same number of reactants and the same number of products execute the same instructions with different values. The revised version groups reactions based on the number of reactants and products, which helps vectorization. Similarly, the dynamical core is annotated to hint the compiler for vectorization. The improvement on the dynamical core is about 1.75x because its performance is mainly limited by the memory bandwidth as predicted by ExaSAT.

In addition, we would like to leverage the lessons learned through ExaSAT for the development of a programming model for combustion codes. The new programming model will focus on data structure support for tiling optimizations for data locality and the use of functional semantics to help the runtime reason about data flow and memory use. The goal is to tune the aggressiveness of tiling and fusion optimizations on a given architecture and minimize data movement.

6.4 Future Work

An area of future work will be to expand the framework’s functionality to cover a broader range of applications besides structured grid problems. For example, we have interest in studying dense linear algebra and N-body problems, which can be statically analyzable. However, the analysis by the compiler and cache model in ExaSAT must be extended to cover their reuse patterns.

One of the current limitations of the framework is how it handles conditionals Vera and Xue (2002). Conditionals come in different flavors and each flavor needs to be handled differently. In the codes we analyzed, the branches contain the same number of memory accesses and only the values assigned to the variables are different. Thus, we only need to analyze one of the if-clauses. When conditionals lead to thread divergence (gaps in the iteration space), we would like to be able compute the data

movement by weighting each branch. If the branches introduce workload imbalance, the model can be parameterized by a branch-taken probability computed from sample runs or provided by the user.

Although we have made substantial progress in identifying several hardware design trade-offs, there are still a number of co-design questions that remain to be answered. We formulated plans for comparing analytic model estimates with dynamic analysis and architectural simulators to obtain more accurate results. Some of these plans include more detailed core model, on-chip network, NVRAM power modeling and network job placement strategies. Another exascale co-design challenge that we have already started evaluating is whether the software or hardware should take responsibility for fault tolerance. Hardware-managed resilience mechanisms increase the overall system cost and power consumption. We are extending our analysis of data access patterns to compute data movement requirements of different checkpoint schemes for software-managed resilience. Finally, given that our methodology allows us to address hardware requirements for the SMC combustion code, we would like to extend the ExaSAT framework to examine the requirements for adaptive mesh refinement codes, such as the Low Mach number Combustion code (Day and Bell, 2000).

7 Conclusions

We developed the ExaSAT framework to rapidly evaluate exascale proxy applications and accelerate the iterative co-design process. ExaSAT complements more detailed architectural simulation tools through rapid generation of abstract analytic models. It is our belief that analytic models are essential to quickly identify the most productive areas for exploring a complicated multi-dimensional design space, including both hardware and software optimizations. ExaSAT parameterizes *both* the machine model and software optimizations to conduct a sensitivity analysis to guide the co-design process. We demonstrated ExaSAT’s ability to perform end-to-end analysis on a combustion proxy application (SMC). The SMC results show substantial opportunities to reduce memory bandwidth requirements by increasing chip area for more registers and on-chip memory. Our analysis illustrates to hardware and software designers the need for higher memory bandwidth and more aggressive software optimizations to reduce data movement. This information can be combined with architectural simulations to understand how our design recommendations interplay with the energy costs of feasible implementa-

tions. Future work will expand the scope of analysis to a wider range of applications and improve the coupling of analytic models with architectural simulation environments.

Acknowledgements

Authors would like to thank Weishen Mead and Matthew Cordery for their contribution to the PinTool validation. All authors from Lawrence Berkeley National Laboratory were supported by the Office of Advanced Scientific Computing Research in the Department of Energy Office of Science under contract number DE-AC02-05CH11231. This work is part of the DOE Center for Exascale Simulation of Combustion in Turbulence (ExaCT) and the DOE Co-Design for Exascale (CoDEx) projects.

References

- Balaprakash, P., Buntinas, D., Chan, A., Guha, A., Gupta, R., Narayanan, S. H. K., Chien, A., Hovland, P., and Norris, B. (2013). Exascale workload characterization and architecture implications. In *21st High Performance Computing Symposia*, HPC.
- Binkert, N., Beckmann, B., Black, G., Reinhardt, S. K., Saidi, A., Basu, A., Hestness, J., Hower, D. R., Krishna, T., Sardashti, S., Sen, R., Sewell, K., Shoaib, M., Vaish, N., Hill, M. D., and Wood, D. A. (2011). The gem5 simulator. *SIGARCH Computer Architecture News*, 39(2):1–7.
- Carrington, L., Snively, A., Gao, X., and Wolter, N. (2003). A performance prediction framework for scientific applications. In *ICCS Workshop on Performance Modeling and Analysis*, PMA03, pages 926–935.
- Caulfield, A. M., Coburn, J., Molloy, T., De, A., Akel, A., He, J., Jagatheesan, A., Gupta, R. K., Snively, A., and Swanson, S. (2010). Understanding the impact of emerging non-volatile memories on high-performance, io-intensive computing. In *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC ’10, pages 1–11, Washington, DC, USA. IEEE Computer Society.

- Cesar (2013). CESAR: Center for exascale simulation of advanced reactors. Website. <http://cesar.mcs.anl.gov/>.
- Chan, C., Unat, D., Lijewski, M., Zhang, W., Bell, J., and Shalf, J. (2013). Software design space exploration for exascale combustion co-design. *Proceedings of the International Supercomputing Conference*.
- Chen, J. H., Choudhary, A., de Supinski, B., DeVries, M., Hawkes, E. R., Klasky, S., Liao, W. K., Ma, K. L., Mellor-Crummey, J., Podhorszki, N., Sankaran, R., Shende, S., and Yoo, C. S. (2009). Terascale direct numerical simulations of turbulent combustion using S3D. *Computational Science and Discovery*, 2(1):015001.
- Day, M. S. and Bell, J. B. (2000). Numerical simulation of laminar reacting flows with complex chemistry. *Combustion Theory and Modelling*, 4:535–556.
- Emmett, M., Zhang, W., and Bell, J. (2013). High-order algorithms for compressible reacting flow with complex chemistry. *under revision on Combustion Theory and Modelling*.
- Ern, A. and Giovangigli, V. (1995). Fast and accurate multicomponent transport property evaluation. *Journal of Computational Physics*, 120(1):105 – 116.
- Exact (2013). ExaCT: Center for exascale simulation of combustion in turbulence. Website. <http://exactcodesign.org>.
- ExaSAT XML Specification (2013). Website. <http://crd.lbl.gov/projects/combustion-codesign-2/>.
- Exascale Research PI Meeting (2012). Combustion Co-Design Center:Exascale Simulation of Combustion in Turbulence Application/Proxy Deep Dive. <http://exactcodesign.org/main/wp-content/uploads/ExaCT-Deep-Dive-Intro.pdf>.
- ExMatEx (2013). ExMatEx: Exascale co-design center for materials in extreme environments. Website. <http://exmatex.lanl.gov/>.

- Fast Forward (2013). FastForward: DOE exascale technology acceleration. Website. <https://asc.llnl.gov/fastforward/>.
- Gottlieb, S. and Shu, C. (1998). Total variation diminishing Runge-Kutta schemes. *Mathematics of Computation*, 67(221):73–85.
- Jung, M., Wilson, E. H., Donofrio, D., Shalf, J., and Kandemir, M. T. (2012). Nandflashsim: Intrinsic latency variation aware nand flash memory system modeling and simulation at microarchitecture level. In *IEEE 28th Symposium on Mass Storage Systems and Technologies, MSST 2012, April 16-20, 2012, Asilomar Conference Grounds, Pacific Grove, CA, USA*, pages 1–12.
- Kamakoti, R. and Pantano, C. (2009). High-order narrow stencil finite-difference approximations of second-order derivatives involving variable coefficients. *SIAM Journal on Scientific Computing*, 31(6):4222–4243.
- Kogge, P., Bergman, K., Borkar, S., Campbell, D., Carlson, W., Dally, W., Denneau, M., Franzen, P., Harrod, W., Hill, K., Hiller, J., Karp, S., Keckler, S., Klein, D., Lucas, R., Richards, M., Scarpelli, A., Scott, S., Snavely, A., Sterling, T., Williams, R. S., and Yelick, K. (2008). ExaScale computing study: Technology challenges in achieving exascale systems. Technical report, DARPA.
- Krasnov, A., Schultz, A., Wawrzynek, J., Gibeling, G., and Droz, P.-Y. (2007). Ramp blue: A message-passing manycore system in fpgas. In *Field Programmable Logic and Applications, 2007. FPL 2007. International Conference on*, pages 54–61.
- Lee, B. C. et al. (2009). Architecting phase change memory as a scalable DRAM alternative. *SIGARCH Computer Architecture News*, 37(3):2–13.
- Li, D., Vetter, J. S., Marin, G., McCurdy, C., Cira, C., Liu, Z., and Yu, W. (2012). Identifying opportunities for byte-addressable non-volatile memory in extreme-scale scientific applications. In *Proceedings of the 2012 IEEE 26th International Parallel and Distributed Processing Symposium, IPDPS '12*, pages 945–956, Washington, DC, USA. IEEE Computer Society.
- Luk, C.-K., Cohn, R., Muth, R., Patil, H., Klauser, A., Lowney, G., Wallace, S., Reddi, V. J., and Hazelwood, K. (2005). Pin: building customized program analysis tools with dynamic instrumenta-

- tion. In *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation*, PLDI '05, pages 190–200, New York, NY, USA. ACM.
- Mohiyuddin, M., Murphy, M., Olike, L., Shalf, J., Wawrzynek, J., and Williams, S. (2009). A design methodology for domain-optimized power-efficient supercomputing. In *SC '09: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, pages 1–12, New York, NY, USA. ACM.
- Narayanan, S. H. K., Norris, B., and Hovland, P. D. (2010). Generating performance bounds from source code. *2012 41st International Conference on Parallel Processing Workshops*, 0:197–206.
- Qiu, J. and Shu, C. (2005). Runge-Kutta discontinuous Galerkin method using WENO limiters. *SIAM Journal on Scientific Computing*, 26(3):907–929.
- Quinlan, D. J. et al. (2002). Treating a user-defined parallel library as a domain-specific language. In *Proceedings of the 16th International Parallel and Distributed Processing Symposium*, IPDPS '02, pages 324–. IEEE Computer Society.
- Qureshi, M. K. et al. (2009). Scalable high performance main memory system using phase-change memory technology. *SIGARCH Computer Architecture News*, 37(3):24–33.
- Rivera, G. and Tseng, C.-W. (2000). Tiling optimizations for 3d scientific computations. In *Proceedings of the 2000 ACM/IEEE conference on Supercomputing*, Supercomputing '00, Washington, DC, USA. IEEE Computer Society.
- Rodrigues, A. F., Hemmert, K. S., Barrett, B. W., Kersey, C., Oldfield, R., Weston, M., Risen, R., Cook, J., Rosenfeld, P., CooperBalls, E., and Jacob, B. (2011). The structural simulation toolkit. *SIGMETRICS Performance Evaluation Review*, 38(4):37–42.
- Saulsbury, A., Pong, F., and Nowatzky, A. (1996). Missing the memory wall: The case for processor/memory integration. In *Proceedings of the 23rd Annual International Symposium on Computer Architecture, Philadelphia, PA, USA, May 22-24, 1996*, pages 90–101. ACM.
- Shalf, J. et al. (2010). Exascale computing technology challenges. In *VECPAR*, volume 6449 of *Lecture Notes in Computer Science*, pages 1–25. Springer.

- Shalf, J. et al. (2011). Rethinking hardware-software codesign for exascale systems. *IEEE Computer*, 44(11):22–30.
- Snively, A., Carrington, L., Wolter, N., Labarta, J., Badia, R., and Purkayastha, A. (2002). A framework for performance modeling and prediction. In *Proceedings of the 2002 ACM/IEEE conference on Supercomputing*, Supercomputing ’02, pages 1–17, Los Alamitos, CA, USA. IEEE Computer Society Press.
- Spafford, K. L. and Vetter, J. S. (2012). Aspen: a domain specific language for performance modeling. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, SC ’12, pages 84:1–84:11, Los Alamitos, CA, USA. IEEE Computer Society Press.
- Thoziyoor, S., Muralimanohar, N., Ahn, J. H., and Jouppi, N. P. (2008). CACTI 5.1. Technical Report HPL-2008-20, HP Labs.
- U.S. Energy Flow Trends (2012). Lawrence Berkeley National Laboratory https://flowcharts.llnl.gov/archive.html#energy_archive.
- Vera, X. and Xue, J. (2002). Let’s study whole-program cache behaviour analytically. In *High-Performance Computer Architecture, 2002. Proceedings. Eighth International Symposium on*, pages 175–186.
- Vladimirov, A. (2012). Arithmetics on Intel’s Sandy Bridge and Westmere CPUs: not all FLOPS are created equal. *Colfax International*.
- Wawrzynek, J., Patterson, D., Oskin, M., Lu, S.-L., Kozyrakis, C., Hoe, J. C., Chiou, D., and Asanovic, K. (2007). RAMP: A Research Accelerator for Multiple Processors. *IEEE Micro*, 27(2).
- Williams, S. et al. (2009). Roofline: an insightful visual performance model for multicore architectures. *Communications of the ACM*, 52(4):65–76.
- Woodward, P. R., Jayaraj, J., Lin, P.-H., Yew, P.-C., Knox, M. R., Greensky, J. B. S. G., Nowatski, A., and Stoffels, K. (2010). Boosting the performance of computational fluid dynamics codes for interactive supercomputing. *Procedia CS*, 1(1):2055–2064.