Assessing the Role of Mini-Applications in Predicting Key Performance Characteristics of Scientific and Engineering Applications

R.F. Barrett, P.S. Crozier, D.W. Doerfler, M.A. Heroux, P.T. Lin, H.K. Thornquist, T.G. Trucano and C.T. Vaughan

*Center for Computing Research
Sandia National Laboratories
Albuquerque, NM, USA*

Abstract

Computational science and engineering application programs are typically large, complex, and dynamic, and are often constrained by distribution limitations. As a means of making tractable rapid explorations of scientific and engineering application programs in the context of new, emerging, and future computing architectures, a suite of “miniapps” has been created to serve as proxies for full scale applications. Each miniapp is designed to represent a key performance characteristic that does or is expected to significantly impact the runtime performance of an application program. In this paper we introduce a methodology for assessing the ability of these miniapps to effectively represent these performance issues. We applied this methodology to three miniapps, examining the linkage between them and an application they are intended to represent. Herein we evaluate the fidelity of that linkage. This work represents the initial steps required to begin to answer the question, “Under what conditions does a miniapp represent a key performance characteristic in a full app?”

1. Introduction

Over the past several years computer architectures commonly employed by the computational science and engineering communities have remained relatively stable, with subsequent generations characterized by faster processors, memories, and interconnects. These changes have typically resulted in predictably faster runtimes for application programs. Emerging and expected future architectures, however, are presenting special challenges and opportunities that, if not effectively exploited, could result in slower runtimes. Driven by stagnant clock speeds, memory constraints, and power consumption restrictions [2, 8], architects are exploring designs involving wider vector units, significantly more but less powerful processor cores, increased threading, more complex memory hierarchies, differently balanced node interconnects, etc.

The means for exploiting these capabilities must be investigated within the relevant context of their use. However, application programs targeting these machines are typically large, complex, dynamic, and often constrained by distribution limitations, and typically outlive the computing environments they originally targeted. They may be constructed using hundreds of thousands to millions of source lines of code, written in multiple programming languages and linking in several third-party libraries, and developed over decades by multiple generations of computational scientists. Thus examining and addressing the issues that will enable effective execution on these machines is prohibitive.

As a means of making tractable rapid explorations of scientific and engineering application programs in this context, a suite of “mini-apps” has been created to serve as proxies for full scale applications. Each miniapp is designed to represent a key performance characteristic that does, or is expected to significantly, impact the runtime performance of a scientific or engineering application program.

These miniapps enable rapid exploration of key performance issues that impact a broad set of scientific application programs. Within the Department of Energy (DOE) they are being used to explore the above issues by a broad set of participants, including staff at DOE laboratories, universities, and vendors. Yet how can we be sure that these proxies adequately represent that which they are intended?

Email address: rffbarre@sandia.gov (R.F. Barrett, P.S. Crozier, D.W. Doerfler, M.A. Heroux, P.T. Lin, H.K. Thornquist, T.G. Trucano and C.T. Vaughan)

The key contribution of the work described herein is a methodology, rooted in formal verification and validation (V&V) efforts that have been developed for experimental science, for determining the quality of the miniapp as it pertains to a large, complex application code. We applied this methodology to three miniapps, examining the linkage between them and an application they are intended to represent. Herein we evaluate the fidelity of that linkage. This work represents the initial steps required to begin to answer the question, “Under what conditions does a miniapp represent a key performance characteristic in a full app?”

1.1. Related work

Application proxies have been part of the code developers’ tool kit for many years. The LINPACK benchmark came into existence [10] as what we are now calling a miniapp. Sweep3d [17], sPPM [3], and the NAS benchmarks [4] in some sense may also be viewed in these terms. Large scale proxies, such as LULESH [18], are serving related purposes. The three DOE Office of Advanced Scientific Computing Research (ASCR) Co-Design Centers ¹ (Ex-MatEx, CESAR, and ExaCT) have identified the development of proxy applications as a key component of their efforts. Miniapps, and other kinds of application proxies, are being used as part of machine procurements ². The Mantevo project [15] solidifies the application proxy idea, bringing a community-based focused effort to bear on the wide variety of explorations that application proxies can enable.

The validation methodology presented herein is the first formal means, that we are aware of, for understanding if, and how, an application proxy may be used to represent the behavior of a full application program. This work is strongly informed by techniques developed for experimental validation, as will be discussed in the following sections.

2. Overview of the Mantevo Project

The Mantevo project [15] was motivated by questions arising from the Trilinos project [14]. These questions concerned the direction of some coding implementations targeting emerging and expected future architectures, including multi-core, many-core, and GPU-accelerated high performance computers. The goal was to create a suite of tools that placed important algorithms into an application-relevant context, enabling rapid exploration of issues and options and their mapping to computing platforms.

Mantevo miniapps are designed and developed to be a tool, useful throughout the co-design space [12], enabling agile exploration of a variety of issues that impact performance. Unlike a compact application, which is designed to capture some sort of physics behavior, miniapps are designed to capture some key performance issues in full applications. Unlike a skeleton application, which is designed for only focusing on inter-process communication perhaps involving a “fake” computation, miniapps create a meaningful context in which to explore key performance issues. Miniapps are developed and owned by application code teams. Miniapps are intended to be modified, and thus are generally limited to a few thousand source lines of code (SLOC), allowing for unconstrained modification. Once no longer useful for these purposes, a miniapp will be discarded. Mantevo miniapps are freely available as open source software under an LGPL license.

The miniapps in the first release of the Mantevo project are listed in Table 1. The first miniapp was HPCCG, which formed and solved a sparse linear system of equations. Although it provides an important capability, it was soon realized that in order to provide a stronger tie to applications of interest, the context in which the linear system is formed needed strengthening. The result was miniFE, putting the linear system into the context of an implicit finite element solver. Thus although HPCCG continues to serve an important role, miniFE will be examined in detail for purposes herein. We anticipate that this sort of situation will continue to occur as these miniapps are used in different situations.

¹<http://science.energy.gov/ascr/research/scidac/co-design/>

²<http://www.nersc.gov/systems/trinity-nersc-8-rfp/draft-nersc-8-trinity-benchmarks/>

<i>Miniapp</i>	<i>Description</i>
CloverLeaf	Solves the compressible Euler equations on a Cartesian grid, using an explicit, second-order accurate method.
CoMD	A simple proxy for the computations in a typical molecular dynamics application. The reference implementation mimics that of SPaSM.
HPCCG	Intended to be the best approximation to an unstructured implicit finite element or finite volume application in 800 lines or fewer.
miniFE	A proxy for unstructured implicit finite element codes. It is similar to HPCCG but provides a much more complete vertical covering of the steps in this class of applications.
miniGhost	A difference stencil across a homogenous three dimensional domain, targeting the inter-process communication halo exchange operation.
miniMD	The force computations in a typical molecular dynamics applications. The algorithms and implementation used closely mimics these same operations as performed in LAMMPS.
miniXyce	SPICE-style circuit simulator [24].

Table 1: List of Mantevo miniapps, release 1.0

3. Methodology

Miniapps are designed to provide a predictive capability for some key performance issue in a full application. Ensuring that a miniapp completely fulfills its intent is a difficult and probably an ongoing task. Further, the runtime behavior of a complex scientific application is typically problem dependent, and therefore it is important to understand the different ways that a code can be used and have a means for configuring the miniapp to mimic the important features under consideration. Thus our approach is to build up a “body of evidence” in support of the goals of a miniapp, combining formal verification and validation (V&V) techniques with our knowledge and experience bases.

Verification is the process of determining that a model implementation accurately represents the developer’s conceptual description of the model and the solution to the model. *Validation* is the process of determining the degree to which a model is an accurate representation of the “real world” (in this case the performance characteristics of the “real” application) from the perspective of the intended uses of the model. These terms are so defined by the American Society of Mechanical Engineers (ASME, 2006) and the American Institute of Aeronautics and Astronautics (AIAA, 1998), and this usage has basically been adopted by the United States DOE and Department of Defense (DoD). That is, within the context of the intent of the comparisons of a model with the “real world,” we must verify that the applications (“real world”) and miniapps (“model”) compare well in the performance dimensions of interest. This is our defined means of assessing that the miniapps are accomplishing what they are designed and required to do (or not). All of the work (and possibly art) in this methodology will be in defining a set of comparisons that allow us to draw conclusions of this kind about the miniapps. We must also understand how close these comparisons should be for us to be able to conclude that the miniapps are suitably accurate models of real code performance, or that they aren’t. There will clearly be significant components of judgment embedded in this methodology given the difficult nature of this problem, but the goal is to achieve some minimal level (at least) of objective evidence that informs us constructively about the fidelity of the miniapps. This approach requires extensive knowledge of, and experience developing, executing, profiling, maintaining, and extending multi-scale, multi-physics scientific and engineering application software, targeting highest performance computing platforms. It also requires a strong understanding of the miniapps and their intended use: what they are intended to represent and what they are not intended to represent. We combine this knowledge into a formal verification and validation (V&V) methodology that lets us examine experimental and predicted data.

This methodology adheres to the spirit of experimental validation as described in [25, 26, 27, 33] because validation referents are intended to be representative of the empirical (that is, “real”) performance of the full applications. However, its important to note that this V&V process is executed within the goals of miniapps, and thus we are not concerned with ensuring any sort of V&V in regard to the application goals such as correctness of the algorithms and output. In other words, in considering the performance fidelity of miniapps, we are not addressing questions about the verification and validation of the full application; we are assuming that this has been done or, if not, that this

is an issue that is not relevant to the immediate goals of developing miniapps. The status of this assumption is in fact of interest, but beyond our immediate scope. Our focus is strictly on the computational runtime characteristics. Beyond this issue we also stress that miniapps are not intended to reproduce specific physics and mathematics represented by the full application. To some degree, we therefore have an operating assumption that a valid miniapp can approximate the runtime performance characteristics of a full application to a useful degree without reproducing the mathematics and physics of the full application to a useful degree. This may be an assumption that is worthy of fuller consideration also, but once again it is beyond the scope of our near term priorities.

3.1. Verification

Mantevo miniapps have been configured so that they produce some outcome that is measurable with some level of confidence of correctness. For example, miniFE solves a partial differential equation (PDE) such that the residual norm of the linear system solution is within an acceptable tolerance.

With regard to the application, we necessarily begin with the assumption it meets its V&V requirements. That said, our methodology still includes a strong element of software verification. In particular, because our validation approach is designed to expose differences in the runtime characteristics between the miniapp and application, the differences seen can point to areas where both the miniapp and application should be examined. In some sense, then, this is a code-to-code V&V exercise, whereby lack of agreement between the two can strengthen the (always ongoing) V&V efforts of each. This of course is not a definitive result, since both codes may be similarly incorrect, and therefore somewhat controversial. However, given the assumptions stated above, this sort of information can still provide a useful service if properly understood.

3.2. Validation

For a set of diagnostic runtime performance characteristics or elements, which we loosely refer to as the *performance domain*,

$$\{D\} = D_1, D_2, \dots, D_n, \quad (1)$$

let

$$\{B\} = B_1, B_2, \dots, B_n, \quad (2)$$

be a corresponding set of baseline full application observational referents, (the “validation data”) and let

$$\{A\} = A_1, A_2, \dots, A_n, \quad (3)$$

be a set of corresponding miniapp measurements.

We then consider the difference between the application referents and the miniapp measurements in the performance domain defined by (1) as some kind of mathematical norm, which we will also call a *validation metric*:

$$X_i = \|B_i - A_i\|_i, \forall i. \quad (4)$$

Here, we have suggested that the difference measurement – the norm – might vary for each component of the performance domain. Clearly, this can become extremely complex and examples presented below will help clarify this. A very simple example could be a situation in which a positive number specifies every component of performance domain. In such a case, we then could simply have:

$$X_i = \|B_i - A_i\|_i = |B_i - A_i|, \forall i. \quad (5)$$

But one component of the performance domain might be specified this way, while another component might be as general as a functional or a time series, in which case the norm is far more complex than simply measuring the absolute value of the difference between two numbers. There is also the possibility that stochastic characteristics must be attached to one or more of the performance dimensions, which further complicates the mathematics that might underlie the validation metric. We do assume that all the components of the performance domain can have a norm distance definition attached to them. In something as extremely complicated as overall runtime performance behavior, even this assumption might fail if we had to deal with qualitative factors in performance. We do not see the need for this level of generality for this discussion.

The point of introducing a validation metric that measures the difference between miniapp performance and the application referent is to draw some conclusion about how well the miniapp is reproducing the performance behavior

of the full application. A simple illustration of this logic is as follows. Suppose that the validity of the miniapp was determined by how accurately it reproduces the performance of the full application in the performance domain. Then the differences in Equation (4) provide the means for assessing validity. Thus, given the measured set of values X_i , for $i = 1, \dots, n$, suppose “valid” accuracy can be assessed using a set of threshold accuracies T_i^1, T_i^2 , for $i = 1, \dots, n$. Then assessment of the validation metric information might then be posed as:

$$V_i = \begin{cases} \text{predictive,} & \text{for } T_i^1 \leq X_i \leq T_i^2 \\ \text{caution,} & \text{for } T_i^2 \leq X_i \leq T_i^3 \\ \text{not predictive,} & \text{for } X_i \geq T_i^3 \end{cases} \quad (6)$$

where V_i is a validity statement attached to performance domain dimension i for some thresholds T_i^j , for $j = 1, \dots, 3$.

While Equation (6) looks like a generally useful algorithm for assessment, we caution there is a great deal of overloading inherent in this simple expression. For example, the choice of thresholds could clearly be extremely difficult. The willingness to even evaluate validity based on a relatively direct threshold assessment is open to debate. And, developing the set $V_i, i = 1, \dots, n$ leaves open the issue of how all of this information is combined into a single appraisal of the validity of the miniapp. Nonetheless, this logic is a clear illustration of the kind of ideal thinking that should underlie the validation assessment of miniapps.

Choice of diagnostics defining the performance domain is clearly a challenge. Example diagnostics for the set D include inter-process communication, which can then be further categorized, such as the number and distribution of partners, and the size and frequency of message traffic. Example observational referents B for the full application may be empirically measured, which we denote as B^M . Or the referents may need to be forecast or estimated using expert judgment, perhaps aided by a tool such as SST [31], denoted B^P . Using referents that are not empirically measured for validation is sometimes called *face validation*. While accepted by several communities, including DoD, clearly *face validation* has weaker validation inference associated with it than the use of directly observed empirical referents. We expect that we might have to use a mixture of both in the validation methodology for miniapps. In general, we emphasize that measurements involving miniapps A must be carefully designed and captured. We also emphasize that having confidence in A and B requires having accumulated meaningful verification evidence.

This framework provides direct advantages. First, the input information D, B , and A and are open to challenge and refinement, are mutable and extensible, and thus the role interpretive judgment in the final results of validity assessment is transparent within the context of use. For example, new diagnostics, new or corrected baseline observations, and new or corrected measurements could be added to the model in the service of better assessment. Second, the way the results are computed can be easily subjected to peer-review scrutiny.

The choice of validation metric can significantly influence differences (the validation metric $\| \cdot \|_i$) between validation data B_i and miniapp measurements A_i . For our validation work, we choose metrics based on their ability to emphasize or highlight differences between the application characteristic under consideration and the miniapp measurements.

For the diagnostics used in this study, the metric was usually configured as a relative comparison between the miniapp and the application. However, when available, the comparison is in an absolute sense.

4. Making the link to full applications

We applied our methodology to four miniapps, examining the linkage between them and an application they are intended to represent. Three of the apps/miniapps will be described here. The fourth app/miniapp pair will be briefly discussed in Section 5 (Summary). Herein we evaluate the fidelity of that linkage. This represents the initial step required to begin to answer the question, “Under what conditions does a miniapp represent a key performance characteristic in a full app?” Runtime profiling information is presented in support of well-specified key performance issues for each code, providing some acceptable level of confidence that the miniapp is representative of their relevant computations, in a manner that will enable exploration and experimentation using the miniapp.

MiniFE was developed to examine Krylov subspace linear equations solution methods applied within the context of an implicit finite element method application on unstructured meshes, such as that used in Charon, a semiconductor device simulator. MiniGhost was designed to provide a means for exploring alternatives to the Bulk-Synchronous Parallel programming model with data aggregation inter-process communication strategy, such as that found in CTH, a shock physics code. MiniMD was developed to model the molecular dynamics Lennard-Jones potential, such as that found in LAMMPS.

In the context of these goals, then, it is important to note that the diagnostics are selected with the intent of understanding individual application-miniapp relationships, and therefore a consistent set of diagnostics across the application-miniapp relationships is neither necessary or expected.

4.1. Experimental Platforms

We have developed and applied our methodology in a broad set of computing environments. Herein we focus on two large scale computers, Cielo and Chama, supplementing the explorations with additional computers where needed to focus on some particular issues. These machines are representative of the ways in which applications take advantage of current capabilities, but also provide hints regarding expected future capabilities.

Cielo [9], an instantiation of a Cray XE6, is designed for production capability computing by the National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) Trilabs, i.e. Sandia National Laboratories (SNL), Los Alamos National Laboratory, and Lawrence Livermore National Laboratory.

Composed of AMD Opteron Magny-Cours oct-core processors, connected using a Cray custom interconnect named Gemini, and a light-weight kernel operating system called Compute Node Linux (CNL). The system consists of 8,944 dual socket compute nodes, for a total of 143,104 cores. Each compute node is divided into two four processor core memory regions, called NUMA nodes (illustrated in Figure 1(a)), connected using HyperTransport version 3.

Figure 1: Node Architectures

Nodes are connected using Cray’s Gemini 3-D torus interconnect. A Gemini ASIC supports two compute nodes. The X and Z dimensions use twice as many links as the Y dimension (24 bits and 12 bits respectively) and therefore introduce an asymmetry to the nodes in terms of bandwidth in the torus. This needs to be taken into account when configuring a system in order to balance the bisection bandwidth of each dimensional slice in the torus. Cielo is configured as a $16 \times 12 \times 24$ 3-D torus. Injection bandwidth is limited by the speed of the Opteron to Gemini HyperTransport link, which runs at 4.4 GT/s. Links in the X and Z dimensions have a peak bi-directional bandwidth of 18.75 GB/s, and the Y dimension peaks at 9.375 GB/s.

Chama, constructed by Appro, Inc. (acquired by Cray in 2012), is designed for production capacity computing for ASC Trilab users. The system consists of 1,232 compute nodes, connected by a QLogic InfiniBand fat tree, using 12000 series switches and 7300 series adapters. Each node is composed of two oct-core Intel Xeon E5-2670 Sandy Bridge processors, illustrated in Figure 1(b), for a total of 19,712 cores, running a RHEL operating system.

Several workstations were employed in order to focus in on some particular issues. We used the following:

- dual socket quad-core Intel Nehalem 5560@ 2.8 GHz processors.
- dual socket quad-core Intel Nehalem 5570@ 2.93GHz processors.
- dual socket oct-core AMD Magny-Cours 6136@2.4 GHz
- dual socket 12-core AMD Magny-Cours @2.1 GHz

4.2. A Molecular Dynamics code

LAMMPS (Large-scale Atomic/Molecular Massively Parallel Simulator) is primarily a classical molecular dynamics (MD) code and can be run in serial or parallel using a spatial-decomposition of the simulation domain [28, 29].

LAMMPS allows simulation of soft materials, solid-state materials, and coarse-grained or mesoscopic systems. In general, it is a parallel particle simulator at the atomic, meso, or continuum scale.

The spatial-decomposition method that LAMMPS uses divides the physical domain into three dimensional sub-boxes, one per parallel process. Each process computes forces on atoms in its box using information from nearby processes (illustrated in Figure 2). As they migrate among processes, the atoms carry along their molecular topology. Communication is via a nearest-neighbor six-way stencil. Parallel scaling would be N/P if load is perfectly balanced. Computation scales as N/P , communication scales as $(N/P)^{2/3}$ (for large problems), and memory scales as N/P .

Figure 2: Molecular dynamics computation. Grid represents spatial-decomposition of the physical domain into subdomains for each parallel process. A single process owns the atoms in the central subdomain that shaded dark gray. Data from the atoms in the light gray ghost region are periodically communicated to that process in order to complete the interatomic force calculations.

MiniMD is designed to model the force computations in typical molecular dynamics applications. The algorithms and implementation used closely mimic these same operations as performed in LAMMPS, with both using the Lennard-Jones potential in our tests.

The work involved in the parallel processing implementation of this model can be categorized as:

1. *Collect off-process ghost atom data (“comm”)*: Point-to-point inter-process communication collects off-process ghost atom data onto the owning parallel process.
2. *Construct neighbors list (“neigh”)*: Each atom is assigned to a cell in a three-dimensional bin. Then using an inter-atomic cutoff distance, a neighbor list (array `neigh`) is constructed for each parallel process.
3. *Compute forces (“forces”)*: Calculate the forces acting upon each atom by the other atoms.

4.2.1. Model Abstractions

As stated throughout this paper, miniapps are not intended to capture all aspects of a particular application. Instead, they are designed to represent issues that critically impact the runtime characteristics of large scale application programs. In addition to enabling a stronger focus on a limited set of important issues in a particular application, this approach allows a miniapp to be representative of more than one application, and in some cases, represent a class of applications and algorithms.

In this case, miniMD does not include the complex logic present in the LAMMPS implementation, resulting in shorter runtimes. Lennard-Jones is one particular type of atomic interaction model that can be used in molecular dynamics. In the future, miniMD may be expanded to include more complex atomic interaction models in order to perform an expanded set of comparisons with LAMMPS. But for now, miniMD cannot explore the broad problem space that is available to LAMMPS, which includes a wide array of complex interatomic interaction potentials and a very long list of auxiliary capabilities (see <http://lammps.sandia.gov> for more details about LAMMPS’s capabilities).

That said, LAMMPS(LJ) and miniMD are quite similar, providing a strong candidate for testing our methodology. This is because the LJ problem and miniMD that we’ve chosen to simulate are quite representative of a typical MD simulation that an MD practitioner would run, particularly in terms of the computational load and communication patterns. The spatial decomposition, neighborhood list builds, force calculations, radial cutoff, and the time integration simulation of a condensed matter fluid are all common themes that are found in most parallel MD calculations. Further, the computations and data access patterns of the force calculations stress the computational capabilities and memory systems in ways that are important to computations found in many other science domain areas.

Processor Cores	L(32k)	MD(32k)	L(16k)	MD(16k)	L(8k)	MD(8k)	L(4k)	MD(4k)
	32k Atoms		16k Atoms		8k Atoms		4k Atoms	
Nehalem Workstation								
1	0.4	0.3	0.3	0.00	0.5	0.1	0.2	0.5
2	0.5	0.6	1.1	0.00	1.2	0.8	0.4	0.9
4	1.0	2.9	0.6	0.01	1.7	0.8	0.7	1.0
8	4.3	1.0	1.0	0.11	7.1	10.2	4.8	0.0
Cielo								
1	0.84	0.14	0.83	0.06	0.03	0.01	0.01	0.03
2	0.67	0.13	0.03	0.05	0.07	0.09	4.28	0.05
4	0.05	0.13	0.03	0.14	5.26	0.06	5.92	0.03
8	0.05	0.18	0.11	0.17	1.80	0.09	2.72	0.06
16	0.07	0.04	1.72	0.02	6.54	0.03	0.18	0.04
32	0.42	0.98	1.45	0.18	0.38	0.15	0.67	0.11
64	4.02	1.69	1.97	1.51	0.21	1.58	0.31	0.12
128	0.98	0.90	1.73	0.53	0.04	0.62	1.70	0.27
256	0.97	0.21	0.84	0.65	0.83	0.25	0.03	0.12

Table 2: Standard deviations, as a percentage of the time, for LAMMPS and miniMD experiments (denoted by “L” and “MD”, resp.).

We have chosen to study scaling in a strong sense, where the number of atoms is held constant with increasing core count, instead of weak scaling, where the number of atoms simulated would be proportional to the number of cores used. This choice is intentional and seems more representative of what a typical MD practitioner would want: run a relatively small ensemble of atoms out to as-long-as-possible timescales, rather than simulate a relatively large problem for a short time. Strong scaling is more difficult from a computer science perspective, but more desirable from an MD practitioner’s perspective. We note, however, that both LAMMPS and miniMD can be used to study strong or weak scaling, and that the conclusions we draw from this strong scaling app-to-miniapp comparison would likely hold for a weak scaling comparison as well.

4.2.2. Performance Domain

The strong connection of miniMD to LAMMPS(LJ) provides a straightforward means of applying our methodology. The diagnostics are defined as the time to solution for each computational phase:

- D_1 : Total time
- D_2 : Force calculation time
- D_3 : Time for construction of neighbor list
- D_4 : Time for inter-process communication

Four different problem sets were defined by varying the number of atoms in the simulation (4000, 8000, 16000, and 32000), each using LJ density $\rho^* = 0.8442$, and LJ temperature $T^* = 1.444$, iterated for 1000 time steps, with an LJ time step size of 0.00462.

The performance for this experiment on Cielo and Chama are illustrated in Figures 3 and 4, respectively. These are the minimum times of three trials for each phase. The graphs on the left are direct comparisons, in terms of time. The graphs on the right show the results of our validation methodology. Here we use the normalized form

$$X_i = (B_i - A_i)/B_i, \quad (7)$$

converted to a percentage.

4.2.3. Discussion

Above we’ve shown the minimum times of three trials for each phase. The average standard deviations are shown in Table 2. Results of similar experiments executed on an Intel Nehalem processor are shown in [6].

Figures 3 and 4 show that miniMD closely mimics LAMMPS performance and scaling characteristics in all cases tested. In nearly every case, miniMD total time is slightly less than LAMMPS total time, with miniMD

(a) Total Time

(b) Force Compute Time

(c) Neighborlist Construction Time

(d) Inter-process Communication Time

Figure 3: LAMMPS and miniMD Strong Scaling on Cielo: Time (left column) and Validation Metric (right column).

(a) Total Time

(b) Force Compute Time

(c) Neighborlist Construction Time

(d) Inter-process Communication Time

Figure 4: LAMMPS and miniMD Strong Scaling on Chama: Time (left column) and Validation Metric (right column).

faster by up to almost 25% and LAMMPS faster in a few cases by a few percent. Examination of each code’s respective force calculation inner loops illustrates why miniMD is typically faster: its inner loop logic is simpler. Although miniMD closely mimics the Lennard-Jones algorithm in LAMMPS, it does so within the context of a single, focused goal. MiniMD allows simulation of only a single LJ component with fixed LJ parameters, whereas LAMMPS allows simulation of multi-component LJ materials with user-prescribed LJ parameters. In LAMMPS the algorithm is contained within the goals of a much larger scope, requiring additional complexity so that code may be shared across algorithms, etc. Consequently, LAMMPS’s additional flexibility takes a small toll in terms of overall performance. This effect becomes less pronounced at larger core counts due to its falling cost relative to other parts of the calculation.

For both LAMMPS and miniMD, the neighbor list construction times are consistently about 12% of the cost of the force calculations. Since both of these costs are computation dominated and similar in nature, this ratio remains fairly constant versus core count, code, and platform. Both codes have been tuned to use pair distance cutoffs that necessitate neighbor list updates at about 20 timestep intervals, near the minimum total computational cost. Calculation costs other than those for neighbor list construction, force calculations, and inter-processor communication were typically less than 3% of the total and never more than 8% in all of the cases examined.

In these strong scaling tests, we see that computation cost dominates at low core count and communication cost dominates at high core count, as expected. This can be seen by noticing in Figures 3 and 4 that the force and neighbor list construction times exhibit near perfect scaling (compute time drops in proportion to the number of cores) in these tests whereas the inter-process communication times remain essentially constant with core count. One should also note that the inter-process communication times are negligible at low core count, but dominant at high core count compared with the force and neighbor list construction times. If these tests were extended to even higher core counts, parallel scalability would continue to deteriorate, yielding no advantage to using more cores.

LAMMPS typically outperforms miniMD in terms of inter-process communication time, presumably due to LAMMPS’s more sophisticated and highly optimized communication code. Interestingly, at higher core counts, the relative performance in this regard tightens, demonstrating remarkable similarity in scalability between the two codes on both platforms. Since communication costs dominate at higher core counts, and since LAMMPS and miniMD use very similar communication patterns, we expect good agreement between the codes in overall timings near the scaling limit. Indeed we do see remarkably similar performance between the codes at the highest core counts, with the maximum difference in total performance less than 11%.

MiniMD computes the forces acting upon each atom, and therefore the actions of each atom (potentially) on the others. This coupling of actions is reflected in a coupling of the computation, which has implications with regard to the computational capabilities of current processors. Several versions of MiniMD are now available for download from the Mantevo website (<http://mantevo.org/>), including OpenMP, OpenACC, OpenCL, and Kokkos (with CUDA, pThreads or OpenMP back-ends) versions. Each version uses a different programming model to parallelize miniMD’s force calculations on the node level. MiniMD is serving as a valuable tool for exploring use of these programming models in LAMMPS.

4.3. A semiconductor device simulation code

Charon is a semiconductor device simulation computer program [13, 23] developed at SNL. It is a transport reaction code used to simulate the performance of semiconductor devices under irradiation. Charon employs the drift-diffusion model, which is a coupled system of nonlinear partial differential equations (PDEs). Finite element discretization (stabilized Galerkin formulation) of these equations in space on an unstructured mesh produces a sparse, strongly coupled nonlinear system. These equations are solved using a Newton-Krylov approach, resulting in a large sparse linear systems of the form

$$AM^{-1}(Mx) = b, \quad (8)$$

for $A \in \mathbb{C}^{N \times N}$, x and $b \in \mathbb{C}^N$, and some preconditioner M . Using functionality from Trilinos [14], the linear systems are solved using a Krylov subspace method iterative solver. An algebraic multigrid preconditioner, with local incomplete factorization as smoothers significantly improves scaling and performance [23]. An example 2D steady-state drift-diffusion solution is illustrated in Figure 5 for a bipolar junction transistor (BJT).

MiniFE is intended to mimic the linear solver computations for a finite element discretization on unstructured meshes. MiniFE solves the steady-state scalar conduction equation as presented in [30]. It assembles finite element matrices into a global matrix and vector, then solves the linear system using the Conjugate Gradient (CG) method. Each finite element is a hexahedron with 8 vertex-nodes. These equations are solved on a three-dimensional box of hexahedra. The domain dimensions are in terms of elements, so for example, a $2 \times 2 \times 2$ box describes eight elements each of which has eight nodes, so it is a $3 \times 3 \times 3$ node domain (27 nodes). The coordinate origin is at the corner

Figure 5: Charon steady-state drift-diffusion solution for BJT with voltage bias. The base, emitter and collector are located in the upper left corner, upper right corner and bottom edge respectively. The shades of gray represent the electric potential, which highest voltage along the bottom and lowest in the upper left portion of domain.

of the global box where $x = 0, y = 0, z = 0$. The box extends along the positive x-axis, positive y-axis, and the *negative* z-axis. Each node corresponds to a row in the matrix. A global identifier, assigned using coordinates and global box dimensions, adds coding convenience to some aspects of matrix-structure generation and finite-element assembly. It should be noted that although the actual miniFE mesh is a structured mesh, as miniFE is intended to mimic aspects of an unstructured mesh application code, miniFE employs unstructured mesh data structures.

The domain is partitioned using the Recursive Coordinate Bisection (RCB) method, and thus some processors own non-contiguous blocks of global node identifiers. Since it is convenient for matrices and vectors to store contiguously-numbered blocks of rows, global node identifiers are mapped to a separate space of row numbers such that each processor's nodes correspond to a contiguous block of row numbers.

4.3.1. Model abstractions

Charon steady-state drift-diffusion problems (examined herein) have three degrees of freedom (DOF) per mesh node. MiniFE solves a linear scalar equation (one degree of freedom per mesh node). Here MiniFE was configured such that the structure of the linear systems mimics those of Charon. It is to be determined if miniFE can be configured to represent this level of complexity.

Runtime for typical Charon problems is dominated by the Krylov solver. The CG solver in miniFE is intended to be sufficiently realistic to be representative of the performance characteristics of the Krylov solver portion in an application code, e.g. to be representative of the scaling of a single Krylov iteration. As miniFE solves a different set of physics compared with Charon, miniFE is not intended to predict either the number of Krylov iterations required or the number of Newton steps required. This would require inclusion of the same physics as the app, which would be under the scope of a compact application rather than miniapp (nomenclature described in Section 2). For the portion of the studies focused on the matrix solve, we compare a single Krylov iteration for both miniFE and Charon.

Performance of the finite element assembly (FEA) is very sensitive to the choice of algorithm. The FEA algorithm employed in miniFE is significantly different from the one employed in Charon. Therefore, one would not expect the performance characteristics of miniFE FEA to be representative of Charon. But it would be of interest to see whether our methodology yields this negative result.

4.3.2. Performance domain

The runtime performance of Charon is dominated by the Krylov solver. These computations are characterized by irregular memory accesses, including across nodes, which informs the choice of diagnostics:

D_1 : Node memory bandwidth

D_2 : Cache hit-to-miss ratio

D_3 : Weak scaling

Figure 6: Effects of memory speeds on the FEA and solver phases of Charon and miniFE. Times, relative to 1333 MHz, are shown in the left column and the validation metrics are in the right column. For the validation metric plots, the vertical axis scale is from 0 to 0.1 rather than 0 to 1 to more clearly see the values and the 1333Mhz results are not plotted as they are zero by definition.

In addition to the two main target architectures in this study (Cielo and Chama), we include Red Sky and two workstation nodes. Red Sky is similar to Chama in that it is a capacity cluster that is composed of Intel Xeon processors, though the older generation Nehalems. The individual nodes, described below, enable configuration of different memory speeds.

4.3.3. Diagnostic: node memory bandwidth

Memory bandwidth within a multicore processor based node is seen as having a significant impact on the performance of many applications [1, 11], including Charon [22]. A typical means for exploring this issue is to vary the number of processor cores employed on the node and comparing the resulting performance efficiency. Although this does vary the amount of memory bandwidth available to a core, it also varies other resources, such as the amount of shared cache. An effective validation study requires stronger evidence to make this claim. However, the fact that the application and miniapp both encountered this issue strengthens the claim that miniFE is representative of the impact of memory speeds on Charon.

Using a dual-socket quadcore Intel Nehalem 5560 clocked at 2.8 GHz and a dual-socket 8-core AMD Magny-Cours 6136 clocked at 2.4 GHz, experiments were configured to better focus on memory bandwidth. The machines were configured to provide memory speeds of 800 MHz, 1066 MHz, and 1333 MHz.

The validation metric is defined as follows: we began by normalizing the performance of the FEA and solver performance (in terms of time), with the performance of the fastest memory setting, and then compared those results between Charon and miniFE, shown in Equation 9.

$$\left| \frac{t^{\text{charon}}(m_{\text{max}}) - t^{\text{charon}}(m_i)}{t^{\text{charon}}(m_{\text{max}})} - \frac{t^{\text{miniFE}}(m_{\text{max}}) - t^{\text{miniFE}}(m_i)}{t^{\text{miniFE}}(m_{\text{max}})} \right|, \quad \text{for } i = 800, 1066, 1333\text{MHz}, \quad (9)$$

where $t(m_{\text{max}})$ is the time spent in the computation at the highest memory speed (1333 MHz). That is, we made a relative comparison between the app and miniapp of the normalized absolute comparison within the app or miniapp.

Figure 6 shows the normalized absolute comparison within the app or miniapp. The FEA phases for Charon

Figure 7: Cache behavior of the FEA and solver phases of Charon and miniFE. Hit rates are shown in left column and the validation metrics in the right column.

and miniFE are not impacted by the change in bandwidth. The solve phase is impacted by memory speed. A 66% increase in memory speed resulted in a 19% increase in solver performance on the Magny-Cours processor for both Charon and miniFE, and on the Nehalem, a 23% and 28% increase for Charon and miniFE, respectively. The validation metrics, illustrated in Figure 6, shows that miniFE is within 4% of all measures of Charon, leading us to claim that, with regard to on-node memory bandwidth, miniFE is predictive of Charon.

4.3.4. Diagnostic: cache performance

This diagnostic considers cache behavior, again with the separation between the FEA and solver phases, using Nehalem (Red Sky compute node) and Magny-Cours nodes, each with three levels of cache (L3 is shared across cores in a die; one die per socket for Nehalem and two dies per socket for Magny-Cours). The hit rate, defined as the proportion of the number of times the processor finds needed data in a cache with the total number of times it looks for data in that cache, plays a significant role in processor performance. This provides us with an absolute performance, so our validation metric is

$$|H_{app} - H_{miniapp}|, \quad (10)$$

where H is the measured cache hit rate. As the miniapp is not the app, we would not expect the miniapp to be predictive of exact cache hit rates. But for the miniapp to be predictive, it should be predictive of a high or low hit rate.

Results are shown in Figure 7. For the FEA phase, Charon and miniFE show strong use of level 1 cache, with a proportional difference of no more than 3%. However, level 2 and 3 hit rates are significantly different, leading us to claim that the cache performance of FEA in miniFE is not predictive of that for Charon. As mentioned previously, as the FEA algorithm employed in miniFE is significantly different from that employed in Charon, we would have expected different performance behavior. As expected, the methodology shows that in terms of cache performance, miniFE FEA is not predictive of Charon. Recall that for memory bandwidth, the methodology did not yield the negative result. This is because for FEA, cache performance is significantly more important than memory bandwidth. Therefore, miniFE FEA does not appear to be predictive of Charon FEA. For the solver phase, we

Figure 8: Charon and miniFE weak scaling for time per Krylov iteration, on Chama and Cielo. Times, relative to 32 processors (for same platform and for same code and preconditioner) are shown in the left column and the validation metrics are shown in the right column.

believe that cache performance is predictive. Although the thresholds for acceptance for level 2 and 3 are arguably high ($\approx 15\%$) the trends are clear.

Most interesting is that the Charon solver’s surprisingly low level 2 hit rate seen on the Nehalem is also seen with miniFE. Careful investigation into this unexpectedly low cache hit rate revealed that the main cause was due to the level 2 cache hardware prefetcher being turned off (BIOS configuration setting). It is important to note that miniFE solve cache behavior was predictive of Charon solve cache behavior when a suboptimal BIOS setting had been selected for the level 2 hardware prefetcher.

4.3.5. Diagnostic: weak scaling

Target simulations involving Charon require execution on very large parallel processing architectures. Effective runtime performance in these environments requires managing the inter-node sharing of data. Effective algorithmic behavior requires managing the convergence characteristics of the linear systems through proper choice of preconditioning techniques.

Here we examine weak scaling characteristics of Charon and miniFE up to 16k core counts. Diagnostics include the Charon Krylov solver with two preconditioning strategies, an incomplete factorization algorithm with no fill (ILU(0)) and a multilevel/multigrid (ML) algorithm. Results for each are analyzed in comparison to miniFE, which does not employ a preconditioner. The idea is that Krylov solvers perform common computations (e.g. addition and scaling of vectors, inner products, and sparse matrix-vector products). Further, applications typically use a breadth of preconditioners, so our goal is to understand where specificity is required and where it is not necessary.

Performance is illustrated in Figure 8. We have not yet determined an effective means for analytically comparing scaling behavior. Instead we can reason about the curves by first noting that performance is different on different architectures, which we believe to be a function of the difference in preconditioning strategies. Although the difference between Charon and miniFE with multigrid preconditioning is large, this is not reason enough to reject the relationship. Instead, we claim that miniFE is *not* predictive of Charon solver with multigrid because miniFE does not include the sorts of computations found in multigrid. Further, an analysis of the interprocess message passing requirements shows that multigrid sends over 40% more message per core than do the other configurations [21].

Figure 9: CTH and miniGhost computation and communication interaction

The difference between Charon with ILU(0) preconditioning and miniFE is less clear, with reasoning driven from the position in the codesign space. For example, from the perspective of some hardware architects, these two approaches are not predictive. However, from the perspective of an algorithm developer perhaps investigating new programming models, miniFE performance could be reasonably predictive. Even with the Charon Krylov solver being similar to CG in miniFE, there still are substantial differences between the two cases, such as the significant difference in how the domain decomposition is performed. Clearly further investigation into this issue is needed. Therefore we assign this diagnostic a *caution* assessment.

4.3.6. Discussion

The evidence shows that miniFE is representative of the cache and memory bandwidth performance of Charon for the solve phase, but not for the FEA phase (since miniFE FEA is not representative of Charon cache performance). Both of these conclusions are aligned with an understanding of the algorithmic implementations of the FEA and Krylov solvers. The scaling studies for large numbers of compute nodes shows meaningful differences between Charon and miniFE. In particular, not surprisingly, the behavior of the multigrid preconditioned Charon solve is not captured by the unpreconditioned miniFE solve. This is an important consideration for future large-scale simulations which are expected to require such sophisticated preconditioning in order to achieve acceptable convergence.

This work illustrates the importance of the comparison metric. Whereas the time-based diagnostics for LAMMPS and miniMD compelled relative comparisons, the availability of an absolute value for cache performance (100%) elucidated the shared behavior between the miniapp and app, especially important for the L2 cache issue on the Red Sky Nehalem processor. Normalizing the comparison to the fastest memory speed clearly showed the relationship of algorithm to this important architectural capability. The lack of a meaningful metric for the scaling behavior, while a weakness, in this case was overcome by a firm understanding of the algorithmic characteristics, and provided the basis for important discussions by the application team. We intend to work toward a meaningful metric in this case, possibly informed by the work below.

4.4. A Shock Physics code

CTH is a multi-material, large deformation, strong shock wave, solid mechanics code developed at Sandia National Laboratories [16]. CTH has models for multi-phase, elastic viscoplastic, porous and explosive materials, using second-order accurate numerical methods to reduce dispersion and dissipation and produce accurate, efficient results.

Several times within a time step boundary information is aggregated and exchanged with up to six neighbors in the grid of processors. In the BSP/message aggregation (BSPMA) model, data from multiple (logical) memory locations are combined into a user-managed array with other data, then subsequently transmitted to the target process. Illustrated in Figure 9 this work incurs three costs: memory utilization (the message buffers), on-node bandwidth (copies into the buffer), and synchronization (leading up to and including the data transfer).

MiniGhost is designed as a proxy for the boundary exchange functionality (likewise also called ghost- or halo-exchange), embedded into the context of difference stencils. This notion of mapping a continuous problem to discrete space and the inter-process communication requirement induced by spatially decomposing the grid across parallel processes adheres to the bulk-synchronous parallel programming model (BSP [34]), arguably the dominant model for implementing high performance portable parallel processing scientific applications [5].

Figure 10: CTH shaped charge simulation. Time progresses left to right.

We have used two distinct problems for CTH. The meso-scale impact in a confined space problem is computationally well-balanced across the parallel processes. This problem involves 11 materials, inducing the boundary exchange of 75 variables. The shaped charge problem, illustrated in Figure 10, involves four materials, inducing the boundary exchange of 40 variables. (This problem was used in the acceptance testing for Cielo [9].) For the shaped-charge problem these messages average 4.1 MB and for the meso-scale problem these messages average 10.4 MB. Process 0 and a couple of other processors near it for the shaped charge problem have more work since they are the genesis of the explosion and thus have additional work relative to the other processes. The runtime trace shows significantly less waiting time than the other cores.

4.4.1. Model abstractions

CTH is a finite volume code, and thus it has values that are based at cells and at the nodes. Its computations are complex, whereas miniGhost computation is rather simple. Prior to sending a message containing boundary data to a neighbor, a CTH MPI process waits for a message from the target processes, which alerts the sending process that the matching receive is posted. The intent is to avoid unexpected messages, potentially a serious issue given the typically large amount of data transmitted. This is a meaningful approach on some architectures, such as Red Storm [32], but is of no help on most, which include a built-in acknowledgment handshake prior to sending messages. Further, the MPI specification provides functionality managing this. Regardless, our intent is to test the capabilities of the MPI implementation on the target architecture external of some means for adapting to the specific capabilities of a particular implementation. That said, one use of a miniapp is to provide a lower impact means for testing other approaches and ideas.

CTH manages data as sets of two dimensional slices, contained in a single pool of allocated memory. This memory management scheme is a relic of past language constraints. MiniGhost allocates distinct three dimensional arrays, each representing a material.

MiniGhost is intended to capture the behavior of the halo exchange typically employed in finite difference and volume codes. Its ability to do so has been demonstrated, and has contributed to performance improvements for CTH [7].

4.4.2. Performance Domain

MiniGhost is designed to capture the inter-process communication requirements of CTH, which is configured using the MPI everywhere model. The diagnostics are defined to measure the message passing characteristics of these requirements:

- D_1 : Number of communication partners (neighbors)
- D_2 : Number of messages per boundary exchange.
- D_3 : Message volume (bytes per neighbor)
- D_4 : Weak scaling

Two problem sets described above (shaped charge and meso-scale) provide the drivers for these measurements. In order to ensure that miniGhost accurately reflects the inter-process communication behavior, it is important to understand the inter-process communication infrastructure, including physical interconnect, logical to physical process maps, system software, and ultimately to how the application manages its communication requirements.

Figure 11: CTH and miniGhost logical processor mapping

For example, when MPI communication is initiated, buffers are configured for managing message queues, unexpected messages, etc. In the distinct communication and computation phases of the full application, the compute stage touches enough data to ensure that all communication data structures have been flushed from the processor cache hierarchy and must be refetched from main memory upon the initiation of the communication phase. The work in miniGhost must be enough in order for this to occur. Thus for example, a single variable weak scaling problem of dimension $100 \times 100 \times 100$ in 8-byte precision means that an eight MByte variable is operated on, stored into another eight MByte variable. Thus 16 MBytes of memory has been traversed by the processor, effectively flushing a cache of that size.

Alternative methods for moving the data are being explored using miniGhost, and therefore the experimenter must take this into consideration.

4.4.3. Diagnostics: Boundary exchange characteristics

MiniGhost was configured to match the number and relative position of communication partners, the number of variables, and the dimensions of those variables such that the size of the messages involved in the boundary exchange are the same as with CTH. Figure 11(a) illustrates the communication pattern of CTH problems, replicated by miniGhost. Diagnostics D_1 , D_2 , and D_3 were configured and verified to be equivalent between miniGhost and the CTH problem sets under consideration, thus by any rational metric, the difference would be zero. This lays the foundation for a more meaningful exploration of communication characteristics, described below.

4.4.4. Diagnostic: Weak Scaling

CTH provides a typical example of a code team adapting to computing architectures: in order to avoid message latencies and exploit global bandwidth, computation is performed across as many variables as possible before a boundary exchange involving those variables can be consolidated into a single message per neighbor. But in a recently completed broad-based study of Cielo capabilities [20], and reproduced on Chama, the nearest neighbor boundary exchange encountered significant scaling degradation beyond 8,000 processor cores. This issue is predicted by miniGhost, illustrated in Figure 12(a). Scaling performance is shown relative to that of one node, i.e. 16 processor cores.

The problem was traced to the mapping of the parallel processes to the three dimensional torus topology. Neighbors in the x direction required a maximum of one hop and in the y direction a maximum of two hops. But the number of hops across the network (referred to as the *Manhattan distance*) was shown to increase significantly in the z direction, as shown in Table 3. This combined with the very large messages of a typical CTH problem set (e.g. for the “shaped charge” problem, 40 three dimensional state variable arrays generated message lengths of almost 5 MBytes) resulted in poor scaling beginning at 8k processes, a trend that accelerated after 16k processes.

In response, we implemented a means by which the parallel processes could be logically re-ordered to take advantage of the physical locality induced by the communication requirements. In the normal mode, CTH (and miniGhost) assigns blocks of the mesh to cores in a manner which ignores the connectivity of the cores in a node. On Cielo, as with other Cray X-series architectures, cores are numbered consecutively on a node, and this numbering continues on the next node. This process re-ordering strategy, illustrated in Figure 11(b), explored using miniGhost and incorporated into CTH, resulted in a significant improvement in scaling performance. As seen in Figure 12(b),

(a) Relative Scaling

(b) miniGhost inter-process communication time

(c) Validation Metric

Figure 12: Performance of CTH and miniGhost. The graphs on the left are from Cielo, the graphs on the right are from Chama.

Number of MPI ranks	Original Order			Re-ordered		
	X	Y	Z	X	Y	Z
16	0.0	0.0	0.0	0.0	0.0	0.0
32	0.0	0.0	0.0	0.0	0.0	0.0
64	0.0	0.0	0.3	0.0	0.3	0.0
128	0.0	0.0	1.0	0.0	0.5	0.0
256	0.0	0.0	1.0	0.0	0.5	0.3
512	0.0	0.1	2.0	0.0	0.6	0.4
1024	0.0	0.3	2.1	0.2	1.0	0.7
2048	0.0	0.3	2.7	0.3	1.2	1.2
4096	0.0	0.3	3.7	0.3	1.2	1.2
8192	0.0	0.5	5.1	0.2	1.1	2.0
16384	0.0	0.5	4.9	0.2	1.1	2.2
32768	0.0	0.5	5.6	0.2	1.1	2.5
65536	0.0	1.1	10.2	0.2	1.6	2.8
131072	0.0	1.1	10.1	0.2	1.6	3.1

Table 3: CTH and miniGhost average hop counts on Cielo

this improvement is attributable to a significant improvement in communication time in the z direction. The remaining scaling issue has been traced to synchronization requirements of the boundary exchange, an issue requiring further study.

The validation metric is defined as the difference between the scaling normalized to 16 cores, as follows:

$$\frac{t^{scth} - t^{smg}}{t^{scth}}, \text{ for time } t; \quad i = 1, \dots, \text{numpes}, \quad (11)$$

where $t^{scth} = \frac{t_i^{cth}}{t_{16}^{cth}}$ and $t^{smg} = \frac{t_i^{mg}}{t_{16}^{mg}}$. This can be reduced to

$$1 - \left(\frac{t_i^{mg}}{t_{16}^{mg}} \right) \left(\frac{t_{16}^{cth}}{t_i^{cth}} \right) = 1 - t^{smg} t^{-scth} \quad (12)$$

The comparison is shown in Figure 12(c). Two things are interesting to note. First, the process re-ordered versions have stronger agreement. Second, at least for Cielo where higher process scales were possible, that agreement remained longer as processor counts increased. This is not surprising since the the increasing Manhattan distance is expected to more strongly impact a more complex code.

This work provides an example of how validation requires a build up a “body of evidence”, as discussed in Section 3. We originally noted the scaling problem in a study involving miniGhost, reproduced it with CTH, identified the issue using miniGhost as the time spent in communication in the z -direction, and tracked this down to the Manhattan distance. This lead to a re-ordering strategy of the parallel processes which reduced the time spent communicating in the z -direction and an overall improvement in weak scaling. The impact of the re-ordering was measured as a reduction in the Manhattan distance and was subsequently measured as providing a positive impact on the time spent stalled attempting to move data on to the network [7]. Further, we confirmed that the computational workload for CTH and miniGhost remained invariant with process counts and that the process re-ordering had no impact on the time spent in computation. From this, we are confident that miniGhost is generally predictive of the weak scaling behavior of CTH.

4.4.5. Discussion

Three diagnostics, D_1 , D_2 , and D_3 , were defined in order to configure miniGhost to capture the general CTH inter-process communication requirements. This created an application relevant context for exploring CTH’s communication behavior.

An issue with scaling performance, tracked to the way in which processes were physically mapped onto a computer, was identified and effectively addressed. We intend to explore the ability of miniGhost to represent the inter-process communication behavior of similar applications that employ dynamic meshes, i.e. adaptive mesh refinement (AMR). This will provide challenges for these diagnostics, but they may inform a process reordering strategy that would

improve the application’s scaling behavior. If this is not the case, AMR may need to be added to miniGhost. Either way, this will compel the definition of an additional diagnostic involving the Manhattan distance.

5. Summary and future work

Miniapps have been shown to provide a tractable means for rapid exploration of a broad set of issues associated with effectively executing large scale scientific and engineering applications on current, emerging, and future architectures. As such, it is crucial to have data-driven evidence of what the miniapp is able to represent, and just as important, a firm understanding of what the miniapp is not able to represent. We presented a methodology for understanding the relationships between an application and the application proxy (the miniapp) in the manner in which it is intended to represent.

All validation work must be driven by usage requirements, which requires a strong understanding of the system under study. Here, runtime profiling points us to the important performance portions of an application, informing a choice of diagnostics. Configuring the appropriate validation metric is critical to elucidating the differences in between the application and the miniapp. We used a variety of formulae for this purpose, in some cases driven by incorrectly chosen metric, in order to clarify differences. The thresholds for acceptance are the most nebulous part of this process. Unlike experimental validation, where a threshold is specified, the threshold in our methodology is viewed more as a guideline than a hard number. Given this, the methodology, in a real sense, provides a formal framework for discussion, interpretation, and challenge. Therefore the process for accessing the miniapp to application connection must be open, iterative, and refined.

The goal of miniapps is to provide actionable information to the application developer. Three distinct applications provided the means for demonstrating this link. MiniMD, closely aligned with the application (LAMMPS) it is intended to represent, enabled calibration of our approach. MiniGhost provides a means for focusing on the inter-process communication requirements of a finite difference or finite volume based application (here CTH), with the computation serving to create a relevant separation between the transmission of data between the parallel processes. Within certain limits, miniFE provides a means to investigate linear solver performance for an implicit finite element method on an unstructured mesh application code (e.g. Charon; for a more detailed study see [21]).

Several issues for each application were found that will require further study. For example, the behavior of the multigrid preconditioner in Charon must be added to miniFE. Our work with molecular dynamics illustrates the need for algorithms in addition to Lennard-Jones. MiniGhost, executed at very large scale, pointed to an issue that will need to be considered as anticipated improvements to interconnect technologies are realized.

Our ongoing work includes examination of a fourth application, Xyce, a circuit modeling tool [19]. Designed to perform transistor-level simulations, the development team has requirements for simulating extremely large circuits that will require strong performance on parallel computing platforms of up to thousands of processors. Although the miniapp miniXyce has provided valuable insights for the code development team, application of the validation methodology revealed that the miniapp does not adequately represent the behavior necessary to explore the performance impacting characteristics in a manner that would contribute to application goals. That said, we view this negative result as a positive step toward understanding the issues associated with accomplishing these performance goals, and will use the validation methodology as a means of modifying miniXyce to provide this capability. Additional information may be found in [6].

Development of additional miniapps in Mantevo is progressing in several areas: a solid mechanics contact algorithm, adaptive refinement of an Eulerian mesh (AMR), some graph-based applications, a new miniapp for Xyce, and others. This work is being meaningfully driven by our validation methodology, guiding implementation strategies and configurations, and enabling an integrated code development model that provides a shorter feedback loop.

Finally, it is important to note that just as the methodology provides a means for building up a body of evidence for a miniapp’s predictive capabilities, we anticipate refining the methodology as it is exercised in an increasingly broader context. In particular, the choice of metric can significantly alter the apparent outcome. In this work we have purposely used somewhat simplistic measures. However, future work will examine the use of more sophisticated metrics. Validation is a never-ending process of discovery.

Acknowledgments

Support for this work was provided through the Advanced Simulation and Computing (ASC) program funded by U.S. Department of Energy’s National Nuclear Security Agency. This effort was greatly enhanced by interactions

with staff throughout Sandia as well as many external organizations. It is heartening to discover the active interests in this work, supported by broad and deep expertise, of the computational science community.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

- [1] S.R. Alam, R.F. Barrett, J.A. Kuehn, P.C. Roth, and J.S. Vetter. Characterization of Scientific Workloads on Systems with Multi-core Processors. In *IEEE International Symposium on Workload Characterization*, 2006.
- [2] S. Amarasinghe et al. ASCR Programming Challenges for Exascale Computing. In *Report of the 2011 Workshop on Exascale Programming Challenges*, July 2011.
- [3] S.E. Anderson and P.R. Woodward. simplified PPM. Laboratory for Computational Science and Engineering, University of Minnesota, 1998.
- [4] D.H. Bailey, E. Barszcz, J.T. Barton, D.S. Browning, R.L. Carter, D. Dagum, R.A. Fatoohi, P.O. Frederickson, T.A. Lasinski, R.S. Schreiber, H.D. Simon, V. Venkatakrishnan, and S. K. Weeratunga. The NAS Parallel Benchmarks. *International Journal of Supercomputer Applications*, 5(3):63–73, Fall 1991.
- [5] R.F. Barrett, S. Ahern, M.R. Fahey, R. Hartman-Baker, J.K. Horner, S.W. Poole, and R. Sankaran. A Taxonomy of MPI-Oriented Usage Models in Parallelized Scientific Codes. In *The International Conference on Software Engineering Research and Practice*, 2009.
- [6] R.F. Barrett, P.S. Crozier, D.W. Doerfler, S.D. Hammond, M.A. Heroux, P.T. Lin, H.K. Thronquist, T.G. Trucano, and C.T. Vaughan. Summary of Work for ASC L2 Milestone 4465: Characterize the Role of the Mini-Application in Predicting Key Performance Characteristics of Real Applications. Technical Report SAND2012-4667, Sandia National Laboratories, 2012.
- [7] R.F. Barrett, C.T. Vaughan, S.D. Hammond, and D. Roweth. Reducing the Bulk of the Bulk Synchronous Parallel Model. *Parallel Processing Letters*, 23(4), December 2013.
- [8] Keren Bergman et al. ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems Peter Kogge, Editor and Study Lead, 2008.
- [9] D.W. Doerfler, M. Rajan, C. Nuss, C. Wright, and T. Spelce. Application-Driven Acceptance of Cielo, an XE6 Petascale Capability Platform. In *Proc. 53rd Cray User Group Meeting*, 2011.
- [10] J. J. Dongarra, P. Luszczek, and A. Petit. The LINPACK Benchmark: Past, Present and Future. *Concurrency Computat.: Pract. Exper.*, 15:803–820, 2003.
- [11] J.J. Dongarra, D. Gannon, G. Fox, and K. Kennedy. The Impact of Multicore on Computational Science Software. *CTWatchQuarterly*, 3(1), February 2007.
- [12] A. Geist and S. Dosanjh. IESP Exascale Challenge: Co-Design of Architectures and Algorithms. *Int. J. High Perform. Comput. Appl.*, 23:401–402, November 2009.
- [13] G.L. Hennigan, R.J. Hoekstra, J.P. Castro, D.A. Fixel, and J.N. Shadid. Simulation of Neutron Radiation Damage in Silicon Semiconductor Devices. Technical Report SAND2007-7157, Sandia National Laboratories, 2007.
- [14] M. A. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu, T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, H. Thornquist, R. Tuminaro, J. Willenbring, and A. Williams. An Overview of the Trilinos Project. *ACM Transactions on Mathematical Software*, 31:397–423, September 2005.
- [15] M.A. Heroux, D.W. Doerfler, P.S. Crozier, J.M. Willenbring, H.C. Edwards, A. Williams, M. Rajan, E.R. Keiter, H.K. Thornquist, and R.W. Numrich. Improving Performance via Mini-applications. Technical Report SAND2009-5574, Sandia National Laboratories, September 2009.
- [16] E.S. Hertel, Jr., R. L. Bell, M. G. Elrick, A. V. Farnsworth, G. I. Kerley, J. M. Mcglaun, S. V. Petney, S. A. Silling, P. A. Taylor, and L. Yarrington. CTH: A Software Family for Multi-Dimensional Shock Physics Analysis. In *Proceedings, 19th International Symposium on Shock Waves*, pages 377–382, 1993.

- [17] Accelerated Strategic Computing Initiative. The ASCI SWEEP3D Benchmark Code. www.llnl.gov/asci_benchmarks, 1995.
- [18] I. Karlin, J. Keasler, and R. Neely. Lulesh 2.0 updates and changes. Technical Report LLNL-TR-641973, Lawrence Livermore National Laboratory, August 2013.
- [19] E.R. Keiter and others. Parallel Transistor-Level Circuit Simulation. In Peng Li, Luis Miguel Silveira, and Peter Feldmann, editors, *Advanced Simulation and Verification of Electronic and Biological Systems*. Springer, 2011.
- [20] S.M. Kelly et al. Report of Experiments and Evidence for ASC L2 Milestone 4467 - Demonstration of a Legacy Application's Path to Exascale. Technical Report SAND2012-1750, Sandia National Laboratories, 2012.
- [21] P.T. Lin, M.A. Heroux, A.B. Williams, and R.F. Barrett. Assessing a mini-application as a performance proxy for a finite element method engineering application code. *In preparation*.
- [22] P.T. Lin and J.N. Shadid. Towards Large-Scale Multi-Socket, Multicore Parallel Simulations: Performance of an MPI-only Semiconductor Device Simulator. *Journal of Computational Physics*, 229(19):6804–6818, 2010.
- [23] P.T. Lin, J.N. Shadid, M. Sala, R.S. Tuminaro, G.L. Hennigan, and R.J. Hoekstra. Performance of a Parallel Algebraic Multilevel Preconditioner for Stabilized Finite Element Semiconductor Device Modeling. *Journal of Computational Physics*, 228(17):6250–6267, 2009.
- [24] L.W. Nagel and D.O. Pederson. SPICE (Simulation Program with Integrated Circuit Emphasis). Technical Report UCB/ERL M382, EECS Department, University of California, Berkeley, Apr 1973.
- [25] W.L. Oberkampf and C.J. Roy. *Verification and Validation in Scientific Computing*. Cambridge University Press, 2010.
- [26] W.L. Oberkampf and T.G. Trucano. Verification and Validation in Computational Fluid Dynamics. *Progress in Aerospace Sciences*, 38, 2002.
- [27] M. Pilch, T.G. Trucano, J. Moya, G. Froehlich, A. Hodges, and D. Peercy. Guidelines for Sandia ASCI Verification and Validation Plans - Content and Format: Version 2.0. Technical Report SAND2000-3101, Sandia National Laboratories, 2000.
- [28] S. Plimpton. Fast Parallel Algorithms for Short-Range Molecular Dynamics. *Journal of Computational Physics*, 117, 1995.
- [29] S. J. Plimpton, R. Pollock, and M. Stevens. Particle-Mesh Ewald and rRESPA for Parallel Molecular Dynamics Simulations. In *Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing*, March 1997.
- [30] J.N. Reddy and D.K. Gartling. *The Finite Element Method in Heat Transfer and Fluid Dynamics*. CRC Press, 2nd edition, 2001.
- [31] A.F. Rodrigues, R.C. Murphy, P. Kogge, and K.D. Underwood. The structural simulation toolkit: exploring novel architectures. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC'06)*, page 157, New York, NY, USA, 2006. ACM.
- [32] J.L. Tomkins, R. Brightwell, W.J. Camp, S. Dosanjh, S.M. Kelly, P.T. Lin, C.T. Vaughan, J. Levesque, and V. Tippiraju. The Red Storm Architecture and Early Experiences with Multi-core Processors. *International Journal of Distributed Systems and Technologies (IJDST)*, 1(2):74–93, May 2010.
- [33] T.G. Trucano, M. Pilch, and W.L. Oberkampf. General Concepts for Experimental Validation of ASCI Code Applications. Technical Report SAND2002-0341, Sandia National Laboratories, 2002.
- [34] L.G. Valiant. A Bridging Model for Parallel Computation. *Commun. ACM*, 33:103–111, August 1990.

Richard Barrett is a Principal Member of the Technical Staff in the Extreme-scale computing group at Sandia National Laboratories. As a member of the Scalable Architectures department in the Center for Computing Research, he leads the Application Performance Modeling and Analysis Team (PMAT), whose goals are to understand and characterize application performance on key HPC platforms that are currently deployed, and to predict performance on future platforms using mathematical modeling methods and techniques and an empirical knowledge base.

Douglas Doerfler is a Distinguished Member of Technical Staff at Sandia National Laboratories. Doug is the ACES Cielo Architect and his research interests include high-performance computer architectures and performance analysis.

Michael Heroux is a Distinguished Member of the Technical Staff at Sandia National Laboratories, working on new algorithm development, and robust parallel implementation of solver components for problems of interest to Sandia and the broader scientific and engineering community. He leads development of the Trilinos Project, an effort to provide state of the art solution methods in a state of the art software framework. Trilinos is a 2004 R&D 100 award-winning product, freely available as Open Source and actively developed by dozens of researchers. In addition to Trilinos, Dr. Heroux works on the development of scalable parallel scientific and engineering applications and maintains his interest in the interaction of scientific/engineering applications and high performance computer architectures. He leads the Mantevo Project, which is focused on the development of Open Source, portable mini-applications and mini-drivers for scientific and engineering applications. Dr. Heroux is a telecommuter for Sandia, maintaining an office at home

in rural central Minnesota and at St. John's University where he is Scientist in Residence in the Computer Science Department. He is a member of the Society for Industrial and Applied Mathematics (SIAM) and past chair of the SIAM Activity Group on Supercomputing. He is a Distinguished Member of the Association for Computing Machinery (ACM). He is the Editor-in-Chief for the ACM Transactions on Mathematical Software, Subject Area Editor for the Journal on Parallel and Distributed Computing and Associate Editor for the SIAM Journal on Scientific Computing.

Paul Lin received his PhD in Mechanical and Aerospace Engineering from Princeton University. He is currently a member of the technical staff at Sandia National Laboratories. His research interests include high performance computing, scalable linear solvers, preconditioners, multigrid methods, and algorithms for multicore processors.

Timothy G. Trucano is currently a Senior Scientist in the technical staff at Sandia National Laboratories in the Optimization and Uncertainty Estimation Department. He received his Ph.D. in mathematical physics from the University of New Mexico in 1980. He began work at Sandia in August of that year. For the first 15 years of his career at Sandia, his focused on research, development, and applications of computational shock wave physics. Over this period of time Trucano worked on problems of shock hydrodynamics, equation of state and constitutive modeling, hypervelocity impact phenomena, nuclear weapons effects, and radiation-

hydrodynamics in high energy density physics applications. Beginning in 1995, his work began to focus on issues of computational science verification and validation (V&V), as well as the application of uncertainty quantification. Currently, he works on technical and programmatic challenges for the NNSA Advanced Simulation and Computing (ASC) V&V Program, as well as research in the application of uncertainty quantification in computational prediction and the accompanying decision environment. For the past five years, Trucano has also been working on the role of V&V and uncertainty quantification in the application of complex social models to national security problems.