

# TADSim: Discrete Event-Based Performance Prediction for Temperature-Accelerated Dynamics

SUSAN M. MNISZEWSKI, CHRISTOPH JUNGHANS, ARTHUR F. VOTER, DANNY PEREZ, and STEPHAN J. EIDENBENZ, Los Alamos National Laboratory

Next-generation high-performance computing will require more scalable and flexible performance prediction tools to evaluate software–hardware co-design choices relevant to scientific applications and hardware architectures. We present a new class of tools called *application simulators*—parameterized fast-running proxies of large-scale scientific applications using parallel discrete event simulation. Parameterized choices for the algorithmic method and hardware options provide a rich space for design exploration and allow us to quickly find well-performing software–hardware combinations. We demonstrate our approach with a TADSim simulator that models the temperature-accelerated dynamics (TAD) method, an algorithmically complex and parameter-rich member of the accelerated molecular dynamics (AMD) family of molecular dynamics methods. The essence of the TAD application is captured without the computational expense and resource usage of the full code. We accomplish this by identifying the time-intensive elements, quantifying algorithm steps in terms of those elements, abstracting them out, and replacing them by the passage of time. We use TADSim to quickly characterize the runtime performance and algorithmic behavior for the otherwise long-running simulation code. We extend TADSim to model algorithm extensions, such as speculative spawning of the compute-bound stages, and predict performance improvements without having to implement such a method. Validation against the actual TAD code shows close agreement for the evolution of an example physical system, a silver surface. Focused parameter scans have allowed us to study algorithm parameter choices over far more scenarios than would be possible with the actual simulation. This has led to interesting performance-related insights and suggested extensions.

Categories and Subject Descriptors: I.6 [**Computing Methodologies**]: Simulation and Modeling

General Terms: Discrete-Event Simulation, Distributed Simulation, Performance

Additional Key Words and Phrases: Accelerated molecular dynamics, temperature-accelerated dynamics

## ACM Reference Format:

Susan M. Mniszewski, Christoph Junghans, Arthur F. Voter, Danny Perez, and Stephan J. Eidenbenz. 2015. TADSim: Discrete event-based performance prediction for temperature-accelerated dynamics. *ACM Trans. Model. Comput. Simul.* 25, 3, Article 15 (April 2015), 26 pages.  
DOI: <http://dx.doi.org/10.1145/2699715>

---

This research has been funded by the Los Alamos National Laboratory (LANL) under the Laboratory Directed Research and Development (LDRD) program. C. J. acknowledges funding by a Los Alamos National Laboratory Director's fellowship. Assigned: Los Alamos Unclassified Report 13-28342. LANL is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. DOE under contract DE-AC52-06NA25396.

Authors' addresses: S. M. Mniszewski, C. Junghans, and S. J. Eidenbenz, Computer, Computational & Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, NM 87545 USA; emails: {smm, junghans, eidenben}@lanl.gov; A. F. Voter and D. Perez, Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545 USA; emails: {afv, danny\_perez}@lanl.gov.

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

2015 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 1049-3301/2015/04-ART15 \$15.00

DOI: <http://dx.doi.org/10.1145/2699715>

## 1. INTRODUCTION

New computer architectures, software stacks, and algorithms will be necessary to further advance scientific research. These changes make challenges related to power consumption, cost of data movement, limited memory, data locality, extreme concurrency, and resilience increasingly pressing, especially as we progress toward exascale. Current architectural paths to exascale include multicore high-end CPUs such as Intel's X86, manycore/embedded simpler low-power cores from embedded systems, and GPU/accelerated highly specialized processors from the graphics/gaming market space such as NVIDIA's Tesla and Intel's Xeon Phi Many Integrated Core (MIC). The high-performance computing (HPC) field has become an adopter of these technologies. The HPC community, driven—among others—by the national security needs of the national laboratories of the U.S. Department of Energy, is involved in “computational co-design” efforts using their traditional software portfolio of physics applications to influence and be influenced by next-generation architectures [U.S. Department of Energy Office of Science 2012; Advanced Simulation and Computing Program 2012]. Computational co-design refers to the close coupling of hardware and software design. Co-design is a very well established set of methodologies in embedded systems [Wolf 1994]; extending these concepts to HPC is a DOE science strategy [U.S. Department of Energy Office of Science 2014].

A view of co-design is that of a performance-guided search and selection from a design space of software and hardware options [Eidenbenz et al. 2012]. For each set of hardware–software choices, we predict the performance using scalable distributed discrete event simulation. The performance measures of interest are runtime metrics such as wall clock time (WCT), speedup, or energy usage. Based on these results, an optimization step—using either a scan of the design spaces or a guided search using optimization techniques such as genetic algorithms, tabu search, or simplex optimization methods—identifies the optimal instance in the search space. Parameterizations define the possible software and hardware options, and each can be represented at a different level of fidelity. For example, when focusing on performance prediction of different algorithmic variants of the software implementation of a method, a coarse model may be sufficient at the hardware level.

The central role of performance prediction via simulation in our co-design approach is the main topic of this work, and here we focus primarily on exploring the software design space. We introduce the concept of a parameterized application simulator. The key stages of an application are modeled as discrete events while abstracting time-intensive parts or kernels of the application. The logic of an application or pseudocode is simulated, including loops, control flow, and termination conditions, similar to a state machine (Figure 1). Designing an application simulator requires identifying the time-intensive elements, quantifying the different algorithm steps in terms of those elements, abstracting them out, and replacing them by the passage of time. A low-fidelity hardware architecture model is defined, considering parameters such as processor speed, core counts available, communications time, and thread overhead. Instrumentation is available for the collection of performance metrics. The software and hardware parameters specified in the simulator define the hardware and software design spaces that we explore. Application simulators allow for fast exploration of application design spaces, testing of new algorithmic extensions before actual implementation, and even testing of proposed hardware architectures that do not yet exist.

Here we use an accelerated molecular dynamics (AMD) method, namely temperature-accelerated dynamics (TAD), to illustrate this application simulation concept. The development of this co-design approach was guided by this algorithmically complex and parameter-rich method.

## 2. THE APPLICATION SIMULATOR CONCEPT

The application simulator approach is most useful for performance analysis of applications that do not have predictable progression prior to runtime. Indeed, performance prediction of typical bulk synchronous parallel (BSP) approaches in most current-day physics simulation codes can often be accomplished by analytical methods. We view PDES as the most appropriate solution for modeling complex applications at the algorithmic level, when cycle-accurate results are not required. As we move toward next-generation HPC, asynchronous programming models are becoming more common, requiring alternative approaches. The use of PDES for performance prediction is then an attractive choice. This approach could be especially powerful in cases where the runtimes of certain sections of the codes are not constant but can vary significantly from one case to the next (e.g., if a nonlinear problem needs to be solved, or when the stop time of some procedure is a random variable). Assessing the performance of various implementation and parallelization strategies would be best approached by direct simulation, as analytical methods would rapidly become tedious.

An application simulator is ideally suited to represent an algorithm that spends most of its runtime performing some atomic unit of work. This can be a frequently used computationally intensive calculation, a data access/movement pattern, or a communication sequence. The algorithm can then be represented as a state diagram of a repeatable loop sequence of stages with a stopping condition, where each stage's runtime is specified in atomic units of work. This allows us to abstract out time-consuming computations, yielding a fast-running DES proxy. Stochastic decision can be included to drive branching, stopping conditions, asynchronous communication, and/or spawning of subtasks. Software, hardware, and algorithm parameters are then defined. The parameters reflect the value of tunable constants or allow for different algorithmic choices. These parameters can be specific or general based on the level of modeling abstraction desired. Instrumentation is added to collect performance, logic-based, and resource usage metrics. Due to the absence of executing work units, precomputed information may be necessary to drive the simulation. Details of the dynamics can be obtained by mining the event progression of a simulation carried out with a standard version of the program.

Currently, building an application simulator is a manual effort, as automated tools are not yet available. Knowledge elicitation from domain experts is required.

## 3. RELATED WORK

The choice of parallel discrete event simulation (PDES) [Banks et al. 2008; Fujimoto 1990; Liu 2010] complements other performance prediction methods that include closed-form analytical models [Bauer et al. 2012; Spafford and Vetter 2012], semianalytical models [Barker et al. 2006, 2009], simulation [Bagrodia et al. 1999; Rodrigues et al. 2011], emulation [Xia et al. 2004; Santhi et al. 2013], and hybrid simulation-emulation [Binkert et al. 2011; Calheiros et al. 2012; Zheng et al. 2004].

We note that the Scalable Simulation Toolkit (SST) [Rodrigues et al. 2011; Hendry and Rodriguez 2012] is a laudable example of PDES use in performance prediction. Unlike our application simulators, SST requires proxy application code as input and produces simulations that run slower than the actual application, which does not allow for exploring large parameter spaces and quick testing of new algorithmic ideas. The faster-running SST/macro coarse-grain simulator runs a skeletonization of an application code for studying the effect of network parameters and topology on performance. A skeleton preserves the control flow of communication code but does not retain the domain algorithm behavior.

The simulation of software simulations is not new. The SIMSIM *metasimulation* has been used for load balancing of parallel and distributed simulation systems [Ewald

et al. 2006] to understand behavior and performance characteristics. Application simulators are used to explore and understand the behavior of systems over many scenarios. A methodology was developed to predict the performance of a distributed simulation of an HLA middleware system prior to implementation [Gianni et al. 2010]. Likewise, the performance of algorithmic extensions can be evaluated prior to implementation using application simulation. The SONSim *second-order simulation* was used to predict performance for computer network simulations [Andelfinger and Hartenstein 2013]. Parameters related to different topologies (e.g., interconnected subnetworks, peer-to-peer networks, and wireless networks) and hardware choices (e.g., Ethernet and Infiniband interconnects) were used to assess whether the simulation would benefit from parallelization. Software and hardware model parameter choices were available to explore the design space for application simulators, and the abstract hardware model allowed for reasonably accurate performance predictions.

In cases where analytical models are too cumbersome, another valuable alternative is experimental algorithmics [McGeoch 2012]. The focus is on empirical methods to aid in building better algorithms and understanding performance given a set of conditions and assumptions [McGeoch 2007]. As an example, this approach has been used to understand cache performance for sorting programs and produce more cache-oblivious ones. Similarly, an application simulator of a physics code can be used to understand the performance under many parameterized scenarios not previously explored.

Simulation was shown to be a viable approach for performance prediction of stochastic algorithms using a machine-learning approach in Jeschke et al. [2011]. Parametrized components or subalgorithms serving as elements of a design space are composed in different combinations to represent the domain code of interest. The use of small benchmark models was shown to provide valuable insight into the algorithm performance.

Another co-design approach to design space exploration is driven by the notion of a mini-application or mini-app from the Mantevo project (mantevo.org). Here, a mini-app is a self-contained proxy for a real scientific code that contains key performance aspects of this type of application. They are written to be amenable to refactoring or change but representative enough to be useful in the scientific problem domain. These open source proxies are available to computer hardware vendors and software stack developers to study and improve application performance. They operate either stand-alone or in simulated environments. In contrast, an application simulator is a virtualization of a scientific code in the form of an event-based simulator. The hardware environment is represented only at the level that is necessary to explore a given domain codebase. Some design choices can be expressed as parameterized options. Both approaches require refactoring or rewrites for exploring different programming models and significant algorithm changes.

#### 4. APPLICATION TO MOLECULAR DYNAMICS

As a first application of this methodology, we target atomistic simulation methods in the molecular dynamics (MD) family. MD is a computer simulation technique for modeling the physical evolution of interacting atoms by numerically solving their equation of motion. To carry out the integration, forces acting on atoms are derived from a potential, a so-called molecular mechanics force field. These calculations, commonly called *force calls*, are the computational core and dominant cycle burner in any MD code. MD is widely used [Rapaport 2004] in materials science [Steinhauser and Hiermaier 2009], chemical physics [Clark 1985], and the modeling of biomolecules [McCammon and Harvey 1988].

Conventional MD allows one to access time scales on the order of hundreds of nanoseconds to microseconds. Efforts to explore longer time scales have led to the development of AMD methods that provide access to time scales on the order of

milliseconds to seconds or more. For many material systems, the dynamical evolution on these longer time scales is characterized by infrequent events, in which the system makes occasional transitions from one state to another (i.e., long periods of uneventful vibrations are punctuated by rare topology changes). The AMD methods exploit this infrequency characteristic to reach longer times.

Three such AMD generalizations have been proposed [Voter et al. 2002; Perez et al. 2009]. These techniques include simulating at a higher temperature to speed up transitions, parallelizing time through running multiple replicas in parallel, or modifying the interatomic potential in a controlled way. In all of these cases, one has to then map the results (in a statistical sense) into the original, unbiased conditions. AMD has been used to investigate a wide range of important phenomena occurring at the atomic scale, such as the velocity dependence of friction during friction force microscopy experiments [Li et al. 2011] and radiation damage annealing processes relevant for nuclear structural materials [Bai et al. 2010], to name only a few. The temperature-raising method, TAD—a complex and algorithmically interesting AMD method—is the subject of the present work.

In our TAD application simulator, TADSim, we exploit the fact that the force call is by far the most computationally expensive part of any (A)MD method, accounting for at least 90% of all cycles. The runtime of a force call depends on the interatomic potential function (the gradient of which is proportional to the forces acting on each atom), the number of atoms, and the number of cores available. The force call is considered to be the atomistic time unit, abstracted out, and represented by the passage of simulation time, enabling fast simulations while allowing for reliable predictions of the application’s runtime. As shown in Figure 1, the TAD algorithm is then expressed as a sequence of stages, each represented by an event, and whose cost is measured in required number of force calls. We focus at a detailed level of the algorithm using specific parameters, whereas the hardware is viewed at a higher level and described using more general parameters.

## 5. TEMPERATURE-ACCELERATED DYNAMICS METHOD

In this section, we will introduce the algorithm and discuss the parameters used for performance prediction. For a detailed discussion of TAD, we refer the reader to Sorensen and Voter [2000]. TAD is an algorithm for reaching long time scales in MD simulations. With few exceptions, direct MD is limited to a maximum simulation time on the order of a single microsecond, due to the requirement that the integration timesteps are short enough ( $\sim 10^{-15}$  s) to resolve the atomic vibrations ( $\sim 10^{-13}$  s). For most materials, the dynamical evolution on these longer time scales is characterized by infrequent events, in which the system makes occasional transitions from one state to another. An individual state is a  $3N$ -dimensional potential basin (where  $N$  is the number of moving atoms), shaped parabolically near the minimum. An example of such an event is the jump of a vacancy in a solid or an atom on a surface, and much more complex events, sometimes involving many atoms, can occur as well. The AMD methods, of which TAD is one, exploit this infrequency characteristic to reach longer times. We chose TAD for this initial study because its rich set of parameters make performance prediction at different settings difficult.

In a direct MD simulation, once an interatomic potential and appropriate boundary conditions have been specified, one simply integrates the classical equations of motion, perhaps modified to take into account coupling to a thermal bath at a desired temperature  $T$ . The MD trajectory, within its microsecond limitation, will then make these occasional transitions from state to state automatically. The system vibrates in this basin many times until a fluctuation causing a large excursion takes the trajectory over a ridge top to an adjacent basin. There is a saddle point associated with this transition (the point on the ridge top with zero gradient of the potential energy and one

negatively curved direction corresponding to the reaction path), and there is typically a large number of adjacent basins in which the transition can take place.

The key to reaching longer time scales for this kind of system is to pick one of the adjacent states, and an associated time, for the next transition. The probability for transitioning to a particular adjacent state is proportional to the *rate constant* for escape to that state. The rate constant to each of these states can be accurately approximated using transition state theory [Marcelin 1915], and if all the adjacent states are known, one escape path can be chosen at random, weighted by its escape rate. Because the first-passage time is exponentially distributed for such rare events,

$$p(t) = ke^{-kt}, \quad (1)$$

where the exponent  $k$  is the sum of the rates over all escape paths out of the state, an appropriate time for the escape can be generated using

$$t_{\text{Random}} = (1/k)\ln(1/r), \quad (2)$$

where  $r$  is a random number distributed uniformly on  $(0,1)$ . This stochastic procedure, known as kinetic Monte Carlo (KMC) [Bortz et al. 1975; Voter 2007], can be used to advance the system from state to state. The list of adjacent states to which the next transition might occur can be generated either by intuition or by a procedure known as adaptive KMC, in which most or all possible saddle points are sought through randomly initiated searches [Henkelman and Jónsson 2001].

Here we focus instead on the AMD approach: we let the trajectory itself find the transition event, as it would in direct MD if we waited long enough, but we modify the dynamics such that the trajectory picks this escape more quickly. The advantage is that only one escape path (or, in the case of TAD, a few escape paths) must be found, releasing us from the burden of trying to find all escape paths. The challenge, though, is to design the modified dynamics in such a way that the relative probabilities of the different escape paths are preserved as accurately as possible. In the TAD method, this acceleration is achieved by raising the temperature of the system while correcting for the tendency for transitions to occur in a different order at high temperature. In the following, we briefly describe TAD, and we define and explain the various TAD parameters that we will vary in the TADSim simulations presented in this article.

In TAD, we advance the system from state to state at a temperature  $T_{\text{Low}}$  using information from simulations at a higher temperature  $T_{\text{High}}$ . Temperature control during these simulations is provided by a Langevin thermostat, such as by using the Langevin-Verlet integrator in Allen and Tildesley [1989], which utilizes a sequence of random numbers. Starting from some configuration point (positions of all the atoms)  $R_{\text{Start}}$  in the initial state (state  $A$ ) of the system, we first thermalize the system by evolving the trajectory for a time  $t_{\text{Therm}}$  so that it loses its memory of the initial condition  $R_{\text{Start}}$  (block 1 in Figure 1). At the end of each thermalization stage (and perhaps multiple times during the thermalization time), we interrupt the trajectory and perform a transition check (block 2 in Figure 1) to verify that the trajectory is still in state  $A$ . If it fails this check (i.e., if it is found to be in a different basin), the trajectory is placed again at some position inside state  $A$ , and the thermalization procedure is started from the beginning. We discuss below how transition checks are performed.

Assuming that the thermalization stage has succeeded, so that we have prepared a properly thermalized trajectory in state  $A$ , the trajectory is continued and official MD time  $t^{\text{High}}$  begins accumulating (block 3 in Figure 1). At regular intervals,  $t_{\text{Block}}$ , the trajectory is interrupted to check for a transition (block 4 in Figure 1), and this integration/detection cycle is repeated until a transition to some state other than  $A$  (e.g., state  $j$ ) is detected.

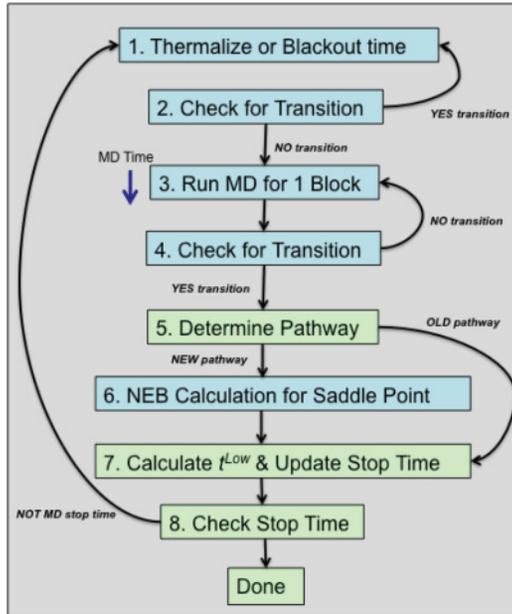


Fig. 1. Serial TAD algorithm as an execution flow of discrete events in TADSim. Blue blocks require force calls and green blocks do not.

Transition checks, which play an important role in TAD, are achieved by using the forces on the atoms to minimize the energy with respect to the geometry  $R$  (e.g., with a steepest-descent or conjugate gradient algorithm) and comparing it to  $R_{\text{Min}}$ , the known minimizer for the current state. If no transition has occurred, it typically takes a few to a few tens of force calls to conclude that the geometry is converging to  $R_{\text{Min}}$ , whereas if the system has made a transition to a new basin, much tighter convergence must be achieved before a transition can be declared, so a larger number of force calls is required. The determination of whether the transition to state  $j$  has been seen before takes place in block 5 of Figure 1.

The time of this first transition to state  $j$ ,  $t_j^{\text{High}}$  is taken to be a random time distributed uniformly on the interval of the most recent block of MD—that is,  $t_j^{\text{High}} = t^{\text{High}} - r t_{\text{Block}}$ , where  $r$  is a random number distributed uniformly on  $[0,1]$ . If this is a transition that has not been seen before, we then initiate a search to find the saddle point that connects state  $A$  with state  $j$  using the nudged elastic band (NEB) method [Jónsson et al. 1998] (block 6 in Figure 1). In brief, the NEB approach relaxes a chain of configurations that connect the initial to the final state, in which each pair of adjacent configurations is connected by a  $3N$ -dimensional spring. At convergence, the highest-energy configuration in this chain gives a good approximation to the saddle-point energy.

A key concept in the TAD method is that each transition event observed in the trajectory at high temperature can be mapped onto a transition event in a corresponding hypothetical trajectory at the low temperature. At  $T_{\text{Low}}$ , the time associated with the transition will be longer, and the order in which the transitions occur will in general not be the same as at  $T_{\text{High}}$ . However, within the harmonic approximation to transition state theory (HTST) [Vineyard 1957], the virtual transitions at  $T_{\text{Low}}$  generated in this way (the escape events and associated times) are in fact indistinguishable from those that would be generated by a long trajectory integrated directly at temperature  $T_{\text{Low}}$ .

In HTST, the escape rate to an adjacent state  $j$  is given by

$$k_j = \nu_{0j} \exp(-E_j/k_B T), \quad (3)$$

where  $\nu_{0j}$  is a temperature-independent pre-exponential factor,  $k_B$  is the Boltzmann constant, and  $E_j$  is the barrier height,  $E_{\text{Saddle},j} - E(R_{\text{Min}})$ . The pre-exponential factor can be computed from the normal-mode vibrational frequencies at the minimum and at the saddle point, although in the TAD method we do not need to do this. Note that in this approximation the only temperature dependence in the rate comes from the barrier height  $E_j$ . Thus, once we have found the saddle point associated with the event, we can compute the ratio of the rate constant at  $T_{\text{High}}$  to the rate constant at  $T_{\text{Low}}$ , and similarly we can compute the appropriately sampled *time* at  $T_{\text{Low}}$  using

$$t_j^{\text{Low}} = t_j^{\text{High}} (k_j^{\text{High}}/k_j^{\text{Low}}) = t_j^{\text{High}} \exp(E_j(1/k_B T_{\text{Low}} - 1/k_B T_{\text{High}})). \quad (4)$$

This gives us a point on our low-temperature time line at position  $t_j^{\text{Low}}$  (block 7 in Figure 1). We now repeat this procedure, beginning with a fresh thermalization in state  $A$  and continuing to accumulate  $t_{\text{High}}$ . Each time we detect a transition, we find the saddle point and use Equation (4) to place a point on the low- $T$  time line.

If we continue this high-temperature trajectory long enough, we will observe one or more events for every possible escape path out of state  $A$ , and each of these events will have an associated time at  $T_{\text{Low}}$  for this particular realization of the dynamics. We can identify the shortest-time event at  $T_{\text{Low}}$  as the transition that would have occurred first at  $T_{\text{Low}}$  (at a time  $t_{\text{Low-Shortest}}$ ), and we can move the system to the state associated with this event.

The second key concept in TAD, the one that allows us to obtain a computational speedup (or *boost*) relative to direct MD, is that we can define a time,  $t_{\text{stop}}$ , at which it is safe to terminate the high-temperature trajectory, knowing that with a desired confidence we have already observed the first low-temperature event. To define this stop time, we make the additional approximation that there is a lower bound  $\nu_{\text{Min}}$  on the pre-exponential factors in the system, such that  $\nu_{0j} \geq \nu_{\text{Min}}$  for all reaction pathways out of  $A$  or any other state the system may visit. We also introduce an uncertainty parameter  $\delta$ , which defines the confidence level  $1 - \delta$  for our assertion that we have found the first event.

We can illustrate the TAD procedure, and derive the stop time, using an Arrhenius-like graphical representation. In a standard Arrhenius graph, one plots the logarithm of the rate against the inverse temperature so that a reaction rate given by Equation (3) corresponds to a downward sloping straight line with slope proportional to  $-E_j/k_B$  and intercept  $\ln(\nu_{0j})$ . We make an analogous plot, replacing the logarithm of the rate with the logarithm of inverse time. On such a plot, as shown in Figure 2, the time progress of the trajectory at  $T_{\text{High}}$  corresponds to nonlinear downward motion on the vertical line at  $1/T_{\text{High}}$ . For each escape event  $j$  at  $T_{\text{High}}$ , the time remapping given by Equation (4) corresponds to extrapolation along a downward-sloping line with slope  $-E_j/k_B$  to find its intersection with the low-temperature time line at  $1/T_{\text{Low}}$ . It is easy to see that a high-barrier event extrapolates to a longer time at  $T_{\text{Low}}$  than a low-barrier event, and this gives rise to the possibility that the events occur in a different order at  $T_{\text{High}}$ . This is the main characteristic for which the TAD procedure must correct: the first event at high temperature is typically not the first event at low temperature.

For any reaction pathway with a rate  $k_j$ , the probability distribution for the first event will be given by an exponential distribution, as in Equation (1). The probability that we will see an escape along this pathway at least once before some time  $t'$  can be found by integrating Equation (1) from zero to  $t'$ , which gives the relationship

$$\text{Prob}(\text{First event time} \leq t') = 1 - \exp(-k_j t'). \quad (5)$$

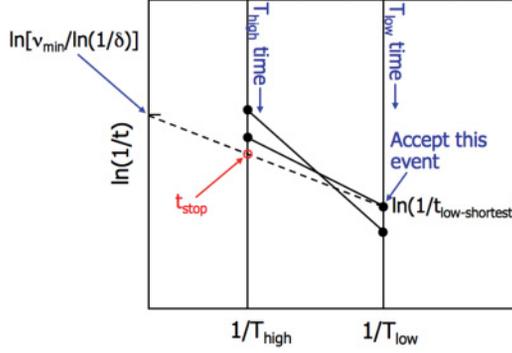


Fig. 2. Pictorial representation of the TAD method. Time moves down the vertical time lines shown. Each attempted transition at the high temperature is extrapolated to the low temperature along a line with a slope given by the negative of the activation energy (see Equation (4)). The MD simulation can be stopped, and the shortest-time event at low temperature can be accepted, when the high-temperature time reaches the intersection with the stop line (dashed), which connects the confidence-modified minimum prefactor on the y axis with the shortest-time event found so far on the low-temperature time line.

If we set this probability to  $1 - \delta$ , and consider a rate with the simple temperature dependence given by Equation (3), we obtain an expression for the time required to be  $(1 - \delta)$ -confident of seeing the event,

$$t_{\text{Conf}} = [v_{0j}/\ln(1/\delta)]\exp(-E_j/k_B T). \quad (6)$$

On our Arrhenius-like plot,  $t_{\text{Conf}}$  corresponds to a line with slope  $-E_j/k_B$  and intercepts  $v_{0j}/\ln(1/\delta)$ . As we move down the vertical time line at a given temperature, when we intersect this line, we have a confidence  $1 - \delta$  that a reactive event along this pathway will have occurred at least once.

We can now define the stop time,  $t_{\text{Stop}}$  (block 7 in Figure 1). Once we have run the high-temperature trajectory long enough to see at least one event, we construct a “stop line,” which connects the time of the current shortest-time event on the low-temperature time line with a point on the y-axis corresponding to the *confidence-modified minimum pre-exponential factor*,

$$v_{\text{Min}}^* = v_{\text{Min}}/\ln(1/\delta), \quad (7)$$

noting that each time a new event is detected, the stop-line definition may change. The stop time is then the intersection of the high-temperature time line with this stop line. When  $t_{\text{High}}$  reaches  $t_{\text{Stop}}$  (block 8 in Figure 1), the probability that any future event at  $T_{\text{High}}$  would create a shorter-time event at  $T_{\text{Low}}$  is lower than  $\delta$ , so we can say with confidence  $1 - \delta$  that it is safe to accept the event at  $t_{\text{Low-Shortest}}$  and move the system to the new state.

This can be understood by considering the scenario in which we proceed with the MD trajectory infinitesimally beyond  $t_{\text{Stop}}$  and encounter an event  $X$  along a new pathway with barrier height  $E_X$ . If  $E_X$  is greater than  $E_{\text{Stop}}$  (where  $E_{\text{Stop}}/k_B$  is the the negative of the slope of the stop line), then the extrapolated time for this new event will be greater than  $t_{\text{Low-Shortest}}$  and there is no problem—event  $X$  would not replace the accepted event. However, if  $E_X$  is less than  $E_{\text{Stop}}$ , then the extrapolated time would be shorter than  $t_{\text{Low-Shortest}}$ . However, by construction, the stop line that we have just crossed is the  $(1 - \delta)$ -confidence line for the reaction pathway corresponding to the worst-case possibility:  $E_X = E_{\text{Stop}}$  and a  $v_X = v_{\text{Min}}$ . A reaction pathway with precisely this barrier and prefactor probably does not exist in the system, but we can say with  $1 - \delta$  confidence that if it did, we would have already seen an event for this pathway. If a slightly different

pathway exists, such as with a higher prefactor, or a lower barrier, then the probability is even greater than  $1 - \delta$  that we would have already observed this pathway. Thus, overall, we can say that the probability is smaller than  $\delta$  that we are accepting the wrong event by terminating the trajectory at  $t_{\text{Stop}}$ .

Assuming that  $\delta$  and  $\nu_{\text{Min}}$  are chosen wisely, the main source of inaccuracy in the TAD method arises from *anharmonicity*, the discrepancy between the actual transition rate for a given pathway and the HTST predicted value for that rate from Equation (3), with its simple temperature dependence. This discrepancy (positive or negative), which is negligible at low  $T$  and typically increases in magnitude with  $T$ , causes an error in the predicted times at  $T_{\text{Low}}$  when Equation (4) is used to convert to the low temperature time for the event.

The TAD method has been implemented in multiple codes, some of which are widely used in the computational physics community [Plimpton 1995].

## 6. TADSIM: MODELING TAD AS A DISCRETE EVENT SIMULATION

TADSim—our parameterized application simulator—models the functionality of the TAD method as described in the previous section through DES of the individual stages of TAD, a flow chart of which is shown in Figure 1. Each stage becomes a discrete event in TADSim, which upon being processed will create future events. The design of TADSim relies on abstracting out the computationally intense parts and replacing them by the passage of time while keeping a realistic execution flow in place. The force call is the most computationally intensive part of TAD and any other MD method. In fact, compute cycle consumption is dominated by force calls in each stage of TAD, and thus the cost of the different stages of TAD is simply proportional to the number of force calls that they consume. This is the key idea behind TADSim.

In the following, we present TADSim details by explaining the parameters available, the execution flow of the TAD simulation, the performance prediction metrics, and the underlying simulation engine.

### 6.1. Design Space Parameters

Parameters that we can vary allow us to explore the hardware and software design spaces. We divide the input parameters into four categories: (1) *Physical System Parameters*, (2) *Method Parameters*, (3) *Architecture Parameters*, and (4) *Simulation Engine Parameters*. Table I lists all available parameters. A description of each category and its associated parameters follow.

*Physical System Parameters* describe the material system that is the target of the TAD simulation.  $N_{\text{Atoms}}$  is the number of atoms simulated.  $T_{\text{Low}}$  is the low temperature at which the physical system is being simulated.  $FcClockCyclesPerAtom$  is the number of clock cycles per atom for each force call. The total number of clock cycles used in a single force call is then  $FcClockCyclesPerAtom \times N_{\text{Atoms}}$ . The actual value of this parameter depends on the complexity of the potential.

A rate catalog of pathway information is precomputed using TAD machinery. This rate catalog gives the characteristics of the various possible transitions that allow the system to exit from the initial state.  $BarrierHeight$  is the energy barrier (the difference in energy between the saddle point and minimum), and  $Prefactor$  is the temperature-independent pre-exponential factor.  $NEBCalls$  is the number of force calls required to find the saddle point using the NEB method.  $AnharmonicityCorrection$  is a parameter used to determine the anharmonicity in the pathway rate when  $T_{\text{High}}$  is higher than the  $AnharmonicityMinTemp$ . The anharmonicity correction factor, which multiplies the rate, is given by  $(1 + (T - AnharmonicityMinTemp) \times (AnharmonicityCorrection - 1) / AnharmonicityMinTemp)$ , where  $AnharmonicityCorrection$  is path dependent. The

Table I. TADSim Simulation Parameters

| Name                                | Description  |
|-------------------------------------|--|
| <u>Physical System Parameters</u>   |  |
| NAtoms                              | Number of atoms in system  |
| $T_{Low}$                           | Low temperature (K)  |
| FcClockCyclesPerAtom                | Number of clock cycles per atom for force call computation         |
| BarrierHeight                       | Barrier height for each pathway in rate catalog (eV)               |
| Prefactor                           | Prefactor for each pathway in rate catalog                         |
| NEBCalls                            | Number of NEB force calls for each pathway in rate catalog         |
| AnharmonicityCorrection             | Correction applied to pathway rate for anharmonicity               |
| AnharmonicityMinTemp                | Minimum temperature (K) for anharmonicity adjustment               |
| FcYesTransition                     | Number of force calls when a transition is detected                |
| FcNoTransition                      | Number of force calls when a transition is NOT detected            |
| FcTransitionWidthFraction           | Fractional width of normal distribution for transition force calls |
| <u>Method Parameters</u>            |  |
| MDTimestep                          | MD Timestep (s)  |
| MinPrefactor                        | Minimum state/pathway prefactor, $\nu_{Min}$                       |
| Delta                               | Uncertainty value, $\delta$  |
| $T_{High}$                          | High temperature (K)   |
| FcBlock                             | Number of force calls per MD block                                 |
| FcThermalize                        | Number of force calls per thermalize                               |
| FcThermalizeCheck                   | Frequency of transition checks during thermalize (force calls)     |
| FcCoreCount                         | Number of cores used for MD computation                            |
| TransCheckSpawnCoreCount            | Number of cores to use when spawning a transition check            |
| NEBBeadSpawnCoreCount               | Number of cores to use per bead when spawning a NEB                |
| NEBBeadCount                        | Number of beads for NEB saddle point calculation                   |
| <u>Architecture Parameters</u>      |  |
| TotalCoreCount                      | Total cores available  |
| CoreClockSpeed                      | Clock speed of cores (GHz)   |
| CommDelay                           | Communications delay during NEB ( $\mu s$ )                        |
| SpawnDelay                          | Delay before and after spawning ( $\mu s$ )                        |
| <u>Simulation Engine Parameters</u> |  |
| TAD_COUNT                           | Number of TAD trials to run  |
| SEED                                | Random seed  |
| TIME_FACTOR                         | Smallest time unit in DES system (fraction of $\mu s$ )            |
| END_TIME                            | End of simulation (based on TIME_FACTOR)                           |

number of force calls required to complete a transition check is represented by a normal distribution, with mean  $FcYesTransition$  and  $FcNoTransition$  for the positive- and negative-resulting transition checks, respectively, and relative width  $FcTransitionWidthFraction$ .  $FcYesTransition$  is much larger because a complete minimization is required when a new state is discovered, whereas a negative result can be determined quickly.  $FcNoTransition$  also has a temperature-dependent factor, given by the same expression as shown earlier for the rate anharmonicity, with  $AnharmonicityCorrection = 6.0$ . These values and characteristics are based on the observed TAD behavior.

*Method Parameters* are software or algorithmic parameters that define the method variation (in our case, a variation of TAD) that we are evaluating and hence describe the software side of the co-design loop. This category contains the most parameters that we will vary in parameter scans in later sections.  $MDTimestep$ ,  $MinPrefactor$ ,  $Delta$ , and  $T_{High}$  have been defined in Section 5.  $FcThermalize$  is the number of force calls required to thermalize the system when the trajectory is replaced in the state after

a transition. Transition checks are performed every `FcBlock` and `FcThermalizeCheck` force calls, during the MD and thermalization stages, respectively. The early detection of a transition during thermalization and/or running MD reduces the amount of wasted MD time, but frequent checks incur additional cost, unless the transition check is spawned off, as we discuss later.

`FcCoreCount` is the number of cores used for running a force computation in parallel. It allows us to test an extension to basic TAD, namely parallelizing a force call. For the modest system sizes that are typically of interest with TAD, the most common parallelization is a force decomposition strategy. The speedup of this approach is limited by synchronization and reduction costs, and hence saturates upon strong scaling. We model the speedup achieved by using `FcCoreCount` cores as  $\sqrt{a \times \text{FcCoreCount}^2 / (a + \text{FcCoreCount}^2)} / \sqrt{a / (a + 1)}$ , where  $a = 100$ . This function initially increases rapidly but saturates at a speedup of 10 for large `FcCoreCount`.

`NEBBeadCount` is the number of beads used in the NEB method (i.e., the number of configurations in the chain). The NEB force calls are executed by each bead in parallel. Additional parallelism is achieved as each bead's force calls can be executed by a separate set of cores.

`TransCheckSpawnCoreCount` and `NEBBeadSpawnCoreCount` specify how many cores can be allocated to transition checks and NEB computations (per bead) respectively, when these computational stages are spawned as separate processes. We describe this in more detail in Section 8.

*Architecture Parameters* characterize the hardware platform. Our hardware model matches the level of abstraction that the software model provides. The architecture parameters describe an architectural model of a large compute cluster consisting of a number of cores (`TotalCoreCount`), each running at the same clock speed (`CoreClockSpeed`). The simulated WCT that a force call takes to execute in TADSim is the product of `CoreClockSpeed`, `NAtoms`, and `FcClockCyclesPerAtom`. Spawning or creating a new process incurs a delay (`SpawnDelay`). If processes on different cores communicate with each other, a communication delay is incurred (`CommDelay`). As we focus mainly on method parameters for TADSim, our hardware model is not very detailed by design.

*Simulation Engine Parameters* control the internals of the underlying simulation engine. These are largely self-explanatory, and we note that the `TIME_FACTOR` is typically set to  $1\mu\text{s}$ .

## 6.2. Execution Flow

TADSim keeps track of multiple time scales simultaneously: simulated WCT for the TAD runtime; MD time at high temperature, collected only when running MD; and MD time at low temperature, as updated after each NEB is calculated. Following the stages of TADSim in Figure 1, thermalization attempts are repeated until no transition is observed during `FcThermalize`. Then MD is run one block at a time (`FcBlock` MD steps) until a transition check indicates that a transition has occurred. The high-temperature MD time is backed up by a random fraction of a block to correct for the overestimation of the real transition time. The transition pathway is determined randomly using the precalculated pathway rates and the algorithm described later. If this pathway was not observed previously, the NEB calculation for the saddle point is performed, at a cost that depends on the nature of the transition. Then, a new putative MD stop time is calculated based on the new TAD low-temperature MD time according to the TAD equations of Section 5. This becomes the official stop time if it is smaller than the previous value. These stages are executed iteratively until the MD stop time is reached.

Table II. TADSim Performance Prediction Metrics

| Name                                   | Description                                   |
|--|---|
| <b>Algorithm Metrics</b>               |   |
| Number of MD force calls               | Total force calls used when running MD        |
| Number of Transition Check force calls | Total force calls used for transition checks  |
| Number of Thermalize force calls       | Total force calls used for running thermalize |
| Number of NEB force calls              | Total force calls used for NEBs               |
| Number of force calls                  | Total force calls                             |
| Number of new pathways                 | Total unique pathways seen                    |
| Number of repeat pathways              | Total pathway replicates seen                 |
| TAD Shortest time                      | Final low-temperature MD time (s)             |
| MD time                                | Total MD time (s)                             |
| MD WCT                                 | Total WCT for MD (s) when running on 1 core   |
| Raw Boost                              | TAD shortest time / MD stop time              |
| Computational Boost                    | MD WCT / WCT                                  |
| Computational Boost Error              | Standard error for Computational Boost        |
| <b>Architecture Metrics</b>            |   |
| Maximum Cores Used                     | Most cores used in parallel during simulation |
| Number of Spawned Transition Checks    | Total transition checks spawned               |
| Number of Spawned NEBs                 | Total NEBs spawned                            |
| Communication Cost                     | Total messages (used when spawning)           |
| WCT                                    | Total runtime (s)                             |

The stochastic nature of TADSim lies in the transition checks. Transition checks during thermalization and running MD involve determining whether a transition has occurred since the last check, and if so, which of the possible transitions occurred and how long the transition check lasted. The probability  $P$  that any transition would occur during the MD block is given by

$$P = 1 - \exp(-k_{\text{Total}} \times \text{MDBlockTime}), \quad (8)$$

where  $\text{MDBlockTime} = \text{FcBlock} \times \text{MDTimestep}$ . If  $P$  exceeds a random number drawn on  $[0,1]$ , then a transition has occurred. Since the probability for each particular transition is proportional to its rate, the index of the transition that occurred can then be obtained by drawing another random number on  $[0,1]$  and identifying the interval of the partial cumulative sum of relative rates where it is located. Finally, the duration of a transition check is drawn from a normal distribution with specified mean (different for positive vs. negative outcomes) and width, as discussed earlier.

### 6.3. Performance Prediction Metrics

TADSim has been instrumented to collect the algorithm and architecture prediction metrics from each TAD trial as shown in Table II. Individual TAD metrics are reported as well as average metrics across all TADs at run completion.

The algorithm metrics provide a means of comparison with the actual TAD simulation runs. The Raw Boost provides the speedup of the TAD method based on MD time alone (i.e., the amount of high-temperature MD time that was required divided by the low-temperature TAD time that was achieved). This does not include the overhead from thermalization and NEBs, and lost time (one half of a block on average) when a transition is detected. This overhead is included in the Computational Boost, defined as the ratio of the WCT that it would take to run the same amount of low-temperature MD time on one core to the predicted TAD WCT. Force call counts are also provided per algorithm stage.

The architecture metrics provide information relevant to the computer architecture being studied, such as the maximum number of cores used in parallel, spawning costs in terms of counts and communication cost, and WCT. The WCT is the predicted runtime of the TAD simulation given the input configuration.

#### 6.4. Simulation Framework

We have used the C++ SimX (formerly known as SimCore) PDES framework [Kroc et al. 2007; Mniszewski 2010] for the implementation of TADSim. SimX is a library for building large-scale distributed-memory, discrete event simulations using the discrete event engine from the Parallel Real-Time Immersive Modeling Environment (PRIME) [Modeling & Networking Systems Research Group, Florida International University 2012] for passing events, event queue maintenance, and synchronization. The important concepts and classes within SimX are Entity, Service, Info, and Profile. Entity is a class that represents a simulation object, such as a Controller or TAD. Service is a class that is used to implement the behavior of an Entity and operates like an event handler. Services are attached to Entities. Info is a class that represents an event that can be scheduled and supplies additional data items and is processed by a Service. Info is passed between Entities (more typically between Services) to trigger an action. Profile is a way of providing runtime specification of default parameter settings for different types of Entities, Services, and Info.

There are two kinds of simulation objects or Entities in TADSim: a TAD Entity and a Controller Entity. A TAD Entity performs the TAD algorithm as shown in Figure 1. Multiple TAD Entities can run in parallel, each representing a separate trial, and no communication is required between them. The simulation can run in a distributed fashion with a Controller Entity on each CPU. The Controller Entity's job is to start the TAD Entities, coordinate transitions, collect prediction metric data from the TAD Entities and report average prediction metrics when done. When running distributed, the Controller Entities operate as a hierarchy for passing on metric results.

### 7. VALIDATION

As a test of our approach, we compare the results of TADSim simulations with real TAD simulations on an actual atomistic system. For simplicity, we choose to investigate a silver adatom (an extra isolated atom) on a silver (100) surface (Figure 3). The system contains 301 atoms, of which 151 are free to move (the bottom layers are held fixed at their equilibrium position). The interatomic potential is of the embedded atom method (EAM) form [Daw and Baskes 1984; Voter 1994], using a parameterization for silver [Voter 1988]. By simulating the system using MD at a high temperature ( $T = 1,200\text{K}$ ) for 25ns, we identified 1,379 possible transitions. The barrier height and pre-exponential factor were determined for each pathway, as needed to compute the HTST rate at any temperature using Equation (3). The NEB calculation for each pathway was deployed in the same way that it would be in an actual TAD simulation, to obtain values for NEBCalls. In this system, the dominant transition pathway (i.e., the one with the lowest barrier of 0.492 electron volts (eV)) is a hopping mechanism that takes the adatom to a nearest-neighbor binding site. There are four of these, as there are four equivalent directions on the Ag(100) surface. At slightly higher energy (0.586 eV), there is an exchange pathway in which the adatom plunges into the surface and pushes a substrate atom up into the next-nearest-neighbor position. All remaining pathways discovered in the MD involve two or more atoms and have barriers in the range of 0.98 eV to 1.65 eV. Anharmonicity corrections were determined, for just the hop and exchange events, by comparing the observed high temperature rates to the HTST rates, giving  $\text{AnharmonicityCorrection} = 1.5$  for the hop and  $\text{AnharmonicityCorrection} = 4.0$

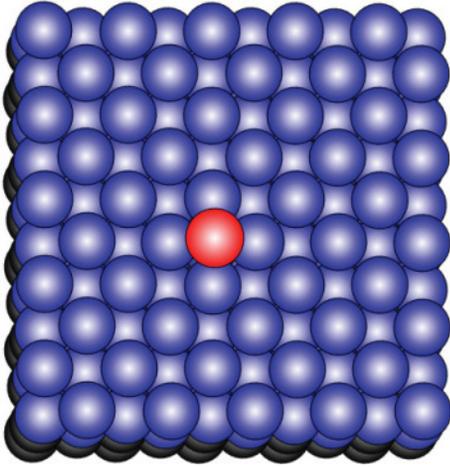


Fig. 3. Atomistic system used for the validation of TADSim, corresponding to an adatom on a silver (100) surface. The moving atoms are blue, the nonmoving atoms are black, and the adatom is shown in red. The adatom can hop to an adjacent binding site or perform a two-atom exchange to a next-nearest-neighbor site. A large number of higher-barrier processes are also available to this system.

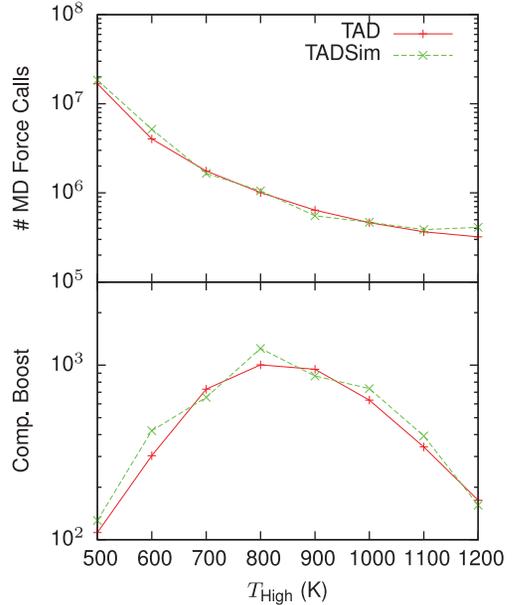


Fig. 4. Comparison of number of MD force calls and Computational Boost for TAD versus TADSim.

for the exchange (i.e., the MD hop rate is 1.5 times faster than the HTST hop rate at  $T = 1,000\text{K}$ , and the exchange rate is four times faster). The resulting catalog of rates was then used by TADSim.

TAD and TADSim were run using the parameter values in the second column of Table III. TAD currently runs on only one core; likewise, the total available cores in TADSim were limited to one. From 60 to 64 trials were run for high temperatures ranging from 500K to 1,200K. The average number of force calls per algorithm stage and the Computational Boost were compared. TAD was run in serial on 2.67GHz Intel<sup>®</sup> Xeon<sup>®</sup> CPU X5650 processors.

The results, shown in Figure 4, demonstrate an excellent agreement between TAD and TADSim results. For example, the number of forces calls taken during MD stages are in excellent agreement over the whole range of high temperatures and properly account for the nearly 100-fold change in this temperature range. We observe a similar agreement in terms of the overall Computational Boost, which is the main metric of interest in this study. Again, TADSim properly reproduced the temperature variation of the boost, including the location of the maximum, and the absolute values are in excellent agreement. We note that the agreement is not as good without the rate anharmonicity corrections.

Incidentally, these results illustrate the challenge in predicting the performance of TAD: whereas the number of force calls decreases monotonously with increasing high temperature (a consequence of the fact that the expected stop time decreases as  $T_{\text{High}}$  increases), the increasingly large overhead from the thermalization, transition checks, and NEBs counteract this gain, leading to an optimal performance at finite temperature. Minimizing the impact of this type of overhead on the runtime of TAD is the key to improving the performance.

Table III. TADSim Parameter Scan Values

| <i>Name</i>                       | <i>Validation</i>                       | <i>Coarse Scan Values</i>                    | <i>Detail Scan Values</i>  |
|-----------------------------------|---|--|--|
| <u>Physical System Parameters</u> |   |  |  |
| NAtoms                            | 151 atoms                               |  |  |
| $T_{low}$                         | 300K                                    | 200K, 300K                                   | 300K   |
| FcClockCyclesPerAtom              | 8,000 cycles                            |  |  |
| BarrierHeight                     | Precomputed per pathway in rate catalog |  |  |
| Prefactor                         | Precomputed per pathway in rate catalog |  |  |
| FcNEB                             | Precomputed per pathway in rate catalog |  |  |
| AnharmonicityCorrection           | Precomputed per pathway in rate catalog |  |  |
| AnharmonicityMinTemp              | 500K                                    |  |  |
| FcYesTransition                   | 500 force calls                         |  |  |
| FcNoTransition                    | 10 force calls                          |  |  |
| FcTransitionWidthFraction         | 0.25                                    |  |  |
| FcThermalize                      | 500 force calls                         |  |  |
| NEBBeadCount                      | 21 beads                                |  |  |
| <u>Method Parameters</u>          |   |  |  |
| MDTimestep                        | 4 fs                                    |  |  |
| MinPrefactor                      | $10^{11} \text{ s}^{-1}$                |  |  |
| Delta                             | 0.01                                    |  |  |
| $T_{high}$                        | 400K to 1,200K<br>in 100K steps         | 400K, 600K, 800K,<br>1,000K, 1,200K          | 800K, 850K, 900K,<br>950K, 1,000K, 1,050K,<br>1,100K, 1,150K, 1,200K |
| FcBlock                           | 500 force calls                         | 1, 2, 10, 50, 100, 400                       | 1, 2, 5, 20, 100   |
| FcThermalizeCheck                 | 500 force calls                         | 1, 2, 10, 50, 100, 400                       | 1, 2, 5, 20, 100   |
| FcCoreCount                       | 1                                       | 1, 16, 64, 512, 4,096 cores                  | 64 cores   |
| TransCheckSpawnCoreCount          | 0                                       | 0, 1, 16, 64, 256 cores                      | 64 cores   |
| NEBBeadSpawnCoreCount             | 0                                       | 0, 1, 16, 64, 256 cores                      | 64 cores   |
| <u>Architecture Parameters</u>    |   |  |  |
| TotalCoreCount                    | 1                                       | 1, 64, 512, 4,096,<br>32,768, $\infty$ cores | $\infty$ cores   |
| CoreClockSpeed                    | 2GHz                                    |  |  |
| CommDelay                         | 1 $\mu$ s                               |  |  |
| SpawnDelay                        | 1,000 $\mu$ s                           |  |  |

These results clearly demonstrate that TADSim reliably predicts performance trends and can be used as a good performance predictor in a systematic co-design approach. This also validates the abstraction of the TAD algorithm that underpins TADSim.

## 8. EXPLORING NEW ALGORITHM EXTENSIONS

As noted previously, mitigating the impact of overhead is key to performance. In that respect, transition checks and NEBs are prime candidates for optimization, given their high cost (up to 500 and 20,000 force calls, respectively, for the validation example).

In usual conditions, most transition checks return with a negative result. It can therefore be beneficial to spawn off a transition check and start running the next MD block, speculating that the transition check will indeed return a negative result. This allows for the concurrent execution of the MD and transition checks. Once a transition is seen, the current MD block is aborted and TADSim rolls back to the point where the MD block that resulted in the positive transition check ended (which can be a few MD blocks in the past, depending on the duration of the check). This is known as a speculative programming model [Tapus and Hickey 2007]. It does require the use of extra cores and start/stop overhead time for the spawning but can lead to a reduction in overall runtime.

Likewise, this same idea can be applied to NEBs. As pathways are seen for the first time, NEBs can be spawned and the next thermalization started, potentially resulting in multiple NEBs executing simultaneously. The stop times are updated as the results of the NEBs become available, and as soon as it can be established that the MD stop time has already been reached by the ongoing MD, unfinished tasks are aborted and the simulation ends. Spawning NEBs can lead to a significant improvement in the runtime.

Speculative spawning requires specifying the number of cores used for a spawned transition check (`TransCheckSpawnCoreCount`) and a spawned NEB (`NEBBeadSpawnCoreCount`). When spawning is attempted, TADSim ensures that a sufficient number of cores are available to complete the request. When not enough cores are available, transition checks are performed serially. In the case of NEBs, the availability of cores is reassessed after a preset delay, until spawning is possible. A delay before and after spawning (`SpawnDelay`) represents the start-up and teardown of lightweight versus heavyweight threads.

## 9. RESULTS

TADSim can be used to predict TAD performance given a set of parameter choices as part of an optimization process or a parameter scan. In the following, the primary performance metric is the Computational Boost, which is the speedup over standard MD on a single core, including overhead. In our co-design context, a set of parameter values defines a variation of TAD on the software side, as well as an architecture on the hardware side. Although a standard parameter optimization procedure, such as a genetic algorithm, could be used for such a task, such a single-path optimization run would not characterize the response surface of the performance function over these hardware and software design spaces very well. In particular, it would not allow us to explicitly study trade-offs between different parameters. However, if the goal is simply to identify optimal parameters without rationalizing the nature of the optimum, an optimization approach is preferable.

We execute parameter scans in a fully factorial fashion over a suitably discretized design space, leveraging access to supercomputing resources. The resulting high-dimensional description of the response surface, which is the dependence of performance on software and hardware parameters, enables us to discover the traces of performance phenomena. However, any factorial design inevitably results in a very significant CPU time consumption. This is addressed by initially running a coarse scan

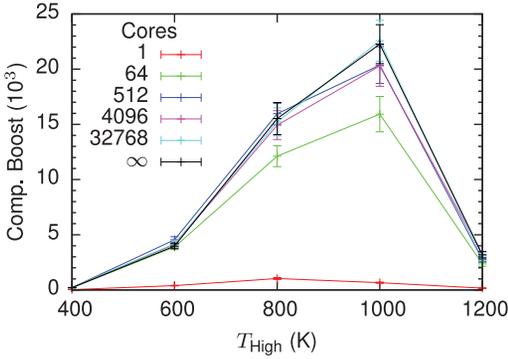


Fig. 5. Maximum Computational Boost as a function of  $T_{\text{High}}$ .

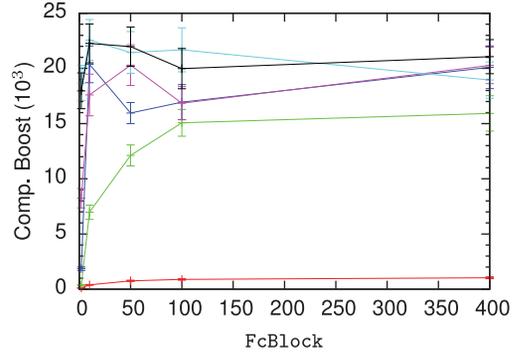


Fig. 6. Maximum Computational Boost as a function of FcBlock (legends are the same as in Figure 5.)

for a discrete set of parameters, followed by a tightly meshed, fully factorial scan around the peak performance point. We further define focused scans on just a few parameters to precisely assess the nature of the performance phenomena.

The evaluation was performed on an HPC network using multiple nodes for simulation. The cluster is composed of nodes with four quad-core AMD Opteron 8354 processors operating at 2.2GHz interconnected using the Infiniband 4X Dual Data Rate DDR network. Single-core DES runs were distributed across multiple nodes using GNU Parallel (<http://www.gnu.org/software/parallel/>).

### 9.1. Parameter Scans to Explore the Performance Response Surface

Here we present results from fully factorial parameter scans of TADSim hardware and software parameters with the values as given in the third column of Table III. Only values that differ from the validation run are shown. The selection of parameters to be varied was chosen by our domain experts. Preliminary runs identified parameters that had little effect on performance, such as CommDelay. This indicates that our variations of TAD are not communication bound. The use of  $\infty$  for TotalCoreCount allowed us to explore extreme or unlimited situations. The small values for the number of force calls between transition checks during thermalization (FcThermalizeCheck) and running MD (FcBlock) represent similarly extreme situations, as they would be very inefficient in serial. Overall, our parameter values generated millions of scenarios. We performed 64 independent runs for each scenario using different random seeds to obtain averaged performance metrics. Partial results of the scan are presented in Figures 5 and 6. In these figures, we vary one parameter,  $T_{\text{High}}$  or FcBlock, and plot the corresponding maximum Computational Boost over the rest of the parameter space. Low temperature ( $T_{\text{Low}}$ ) is set to 300K. Results are shown for different TotalCoreCount limits.

Because the results in Figures 5 and 6 represent scenarios giving the maximum Computational Boost taken from data with statistical error bars, we require an additional procedural step to eliminate bias in the results. For a given parameter constraint (e.g.,  $T_{\text{High}}$  and TotalCoreCount in Figure 5), the optimum scenario chosen from all possible combinations of the remaining parameters will have the largest Computational Boost, not solely due to having the best parameter settings but instead due to a combination of good parameter settings and a favorable (positive) statistical fluctuation. Consequently, this value will be biased in the positive direction. To eliminate this bias, after selecting the maximum-boost scenario, we recompute the Computational Boost for that scenario with a freshly chosen random number seed and plot this value. Each result shown thus represents a lower bound on the true value, with statistical error

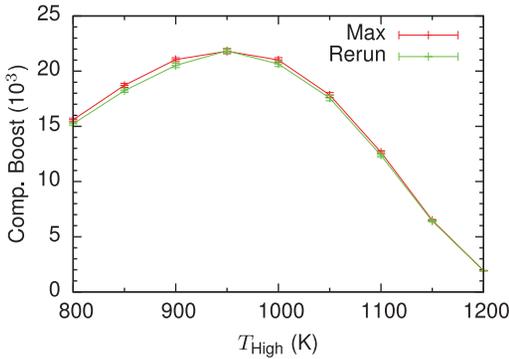


Fig. 7. Maximum Computational Boost versus  $T_{\text{High}}$  in the detailed scan. Optimum performance is achieved at 950K.

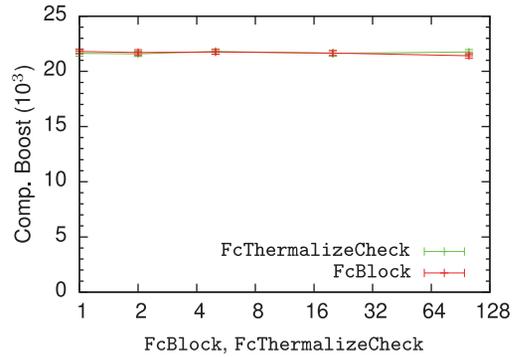


Fig. 8. Maximum Computational Boost versus block size in the thermalization stage and the MD stage for the detailed scan.

bars, where the lower-bound nature arises from two sources: (1) the procedure just described can lead to selection of a less-than-optimal grid point, and (2) the parameter grid has finite resolution.

Figure 5 shows the performance of TAD variations of different maximum core counts (TotalCoreCount) at different high temperatures ( $T_{\text{High}}$ ). Consistently, peak performance is seen at a  $T_{\text{High}}$  of 1,000K for each of the TotalCoreCount values. The corresponding plot for  $T_{\text{Low}} = 200\text{K}$  (not shown) is qualitatively similar but has a less pronounced peak. Even with this coarse scan, we can state that 64 cores are sufficient to reap a majority of the parallel performance possible. Values of TotalCoreCount beyond 64 provide an additional improvement of approximately 40%. We note that the lower boost at 64 cores is limited by a lower number of CPUs available to do MD (FcCoreCount = 16, giving 8.5 times speedup compared to a single core, out of a maximum possible speedup of 10 times) and fewer CPUs for spawning transition checks (TransCheckSpawnCoreCount = 16).

Figure 6 shows the dependence of the performance on FcBlock, the time interval between transition checks. We note again the benefit of allowing additional cores, especially for very frequent transition checks. Interestingly, we see that performance is similar for all values of FcBlock for larger core counts, as all transition checks and NEBs are spawned. Lower core counts result in a mix of spawned and not-spawned transition checks, producing lower performance when FcBlock is small.

The fully factorial scan provides (1) a strong indication of the location of the optimal performance area and (2) uncovers performance behavior phenomena that merit further study. In the following sections, we present the results of a fully factorial tightly meshed scan around the optimum parameter values suggested by the coarse scan and describe a series of low parameter dimension studies to illuminate the performance behavior found in the coarse scan.

## 9.2. Detailed Scan around the Optimum Parameter Values

For the detailed scan, we vary the parameters around the optimal values found in the coarse scan as described in the fourth column of Table III. We applied the same unbiasing procedure as earlier, so again these points represent lower bounds on the exact results, with statistical error bars, although now the bounding and the error bars are much tighter (e.g., see Figure 7, which shows both the maximum points (red) as well as the points resampled with fresh random number seeds (green)). Figure 7 shows the performance results with respect to  $T_{\text{High}}$  when  $T_{\text{Low}} = 300\text{K}$ . We find that for our

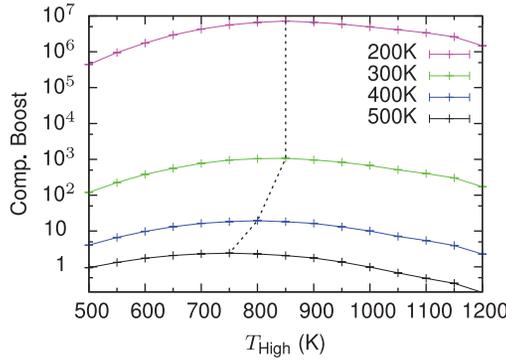


Fig. 9. Computational Boost as a function of  $T_{\text{High}}$  for different  $T_{\text{Low}}$ . The dotted line passes through the optimum  $T_{\text{High}}$  for each  $T_{\text{Low}}$  curve.

silver physical system, the optimum high temperature is 950K. Figure 8 shows the variation of the boost upon varying the time between transition checks during thermalization (FcThermalize) or MD (FcBlock). In this optimization with unlimited cores, any dependence on FcBlock virtually disappears over the range considered ( $\text{FcBlock} \leq 100$ ). This offers the user the opportunity to make more frequent transition checks to improve the physical time-scale resolution, which can be important to the accuracy of the final physical predictions in some cases. However, as discussed previously (also see later discussion), greater total core counts are required for the small values of FcBlock.

### 9.3. Exploring Performance Phenomena: One-dimensional Parameter Scans

The fully factorial coarse scan points to performance phenomena that warrant further study. We focus on a few scenarios in detail, varying only a few parameters, and we organize our findings along observations.

**OBSERVATION 1.** *The optimal high temperature  $T_{\text{High}}$  decreases with increasing  $T_{\text{Low}}$ .*

Our observation of how  $T_{\text{Low}}$  and  $T_{\text{High}}$  relate is at first glance surprising and appears counterintuitive. More in-depth theoretical analysis has confirmed that finding to be accurate. We show in Figure 9 how the optimum high temperature  $T_{\text{High}}$  increases with decreasing  $T_{\text{Low}}$ . The logarithmic y-axis scale is noteworthy, as it illustrates how much more Computational Boost we gain for lower temperatures: more than three orders of magnitude between the low temperatures of 300K and 200K. This is expected given that the exponential decrease of transition rates with decreasing temperature tremendously increases the potential for acceleration. The results also show that the optimum  $T_{\text{High}}$  moves from 850K at a  $T_{\text{Low}}$  of 200K to 750K at a  $T_{\text{Low}}$  of 500K. This can be explained by the following: for a given system, a given setting for  $T_{\text{High}}$  and  $T_{\text{Low}}$  leads to a particular value for the average slope of the “stop line” (dashed line in Figure 2), the line connecting the shortest-time event at  $T_{\text{Low}}$  with  $\nu_{\text{Min}}^*$  on the y-axis. For typical values of  $\nu_{\text{Min}}$  and  $\delta$ , if  $T_{\text{Low}}$  is increased, the average slope of this stop line will decrease in magnitude (become less negative), because the time of the shortest-time event decreases. This weaker slope means that if we consider lowering  $T_{\text{High}}$ , which has the benefit of reducing the average number of NEBs required before the stop time is reached, we can do so with less of a penalty in terms of lost boost. Thus, if  $T_{\text{High}}$  was chosen optimally for the first value of  $T_{\text{Low}}$ , this implies that for a higher  $T_{\text{Low}}$ , the new optimum  $T_{\text{High}}$  is lower. Another consequence of this is that if  $T_{\text{Low}}$  is raised far enough,

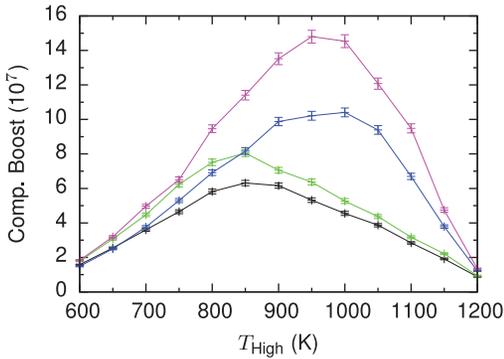


Fig. 10. Computational Boost for no spawning (black), spawning of transition checks (green), spawning of NEBs (blue), and spawning both (red) at  $T_{Low} = 200K$ . Here, 64 cores are used for the MD; 64 cores are also used for each spawned transition check and each NEB bead. (Legend is the same as in Figure 11.)

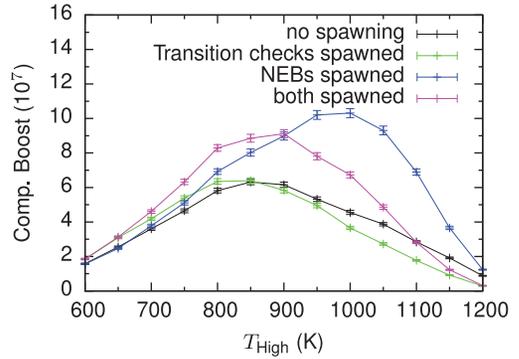


Fig. 11. Computational Boost for no spawning, spawning of transition checks, spawning of NEBs, and spawning both at  $T_{Low} = 200K$ . Here, 64 cores are used for the MD, whereas only one core is used for each spawned transition check and each NEB bead.

the slope of the stop line can actually become positive, at which point the optimum setting for  $T_{High}$  is  $T_{High} = T_{Low}$ , indicating that direct MD will be faster than TAD.

We note that the optimum high temperature in Figure 9 is different (lower) than in our factorial scans because the scenarios used in Figure 9 correspond to  $TotalCoreLimit = 1$ ,  $FcCoreCount = 1$ ,  $FcBlock = 500$ ,  $FcThermalizeCheck = 500$ , and no spawning allowed. In essence, our one-dimensional scan along the  $T_{Low}$  variable was done on a single-core architecture. The absence of spawning accounts for a factor of a few in loss of boost, whereas the serial nature of the force call costs an additional factor of about 10. The fact that the optimum temperature depends on other parameters points to the high-dimensional and complex nature of our response surface. TAD can run at higher optimal  $T_{High}$  values for more complex architectures because the parallelization of force calls and spawning of NEBs and transition checks can partially amortize the cost of processing the larger number of transitions that will occur before the stop time is reached.

**OBSERVATION 2.** *Spawning transition checks and NEBs is advantageous if enough cores are provided.*

Spawning of transition checks and NEBs are two key algorithmic variations for enhancing the basic TAD method. A comparison of TADSim with and without spawning is shown in Figures 10 and 11. The parameters used for these scenarios are as follows:  $T_{Low} = 200K$ ,  $FcBlock = 100$ ,  $FcThermalizeCheck = 100$ ,  $FcCoreCount = 64$ , and  $TotalCoreCount = \infty$ . Figure 10 shows the case where  $TransCheckSpawnCoreCount = NEBBeadSpawnCoreCount = FcCoreCount = 64$ . We see that spawning transition checks alone (green curve) increases the boost by about 25%. Spawning NEBs alone (blue curve) gives a larger effect ( $\sim 60\%$ ), and the optimum  $T_{High}$  increases from 850K to 1,000K, because the spawned NEBs hide most of the work caused by the larger number of attempted transitions at higher  $T_{High}$ . Once NEBs are spawned, spawning transition checks also (red curve) gives a larger effect than spawning transition checks alone. This is because the force-call savings represent a greater fraction of the time once the NEB cost has been eliminated. Note also that the optimum  $T_{High}$  then moves back down a bit, to 950K.

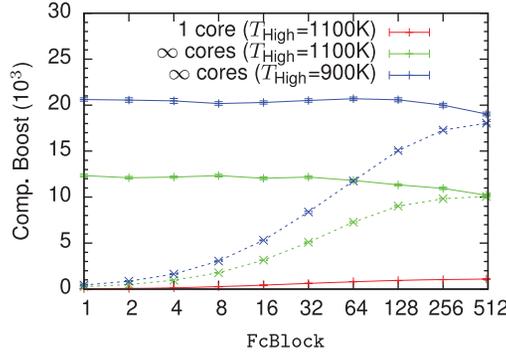


Fig. 12. Computational Boost as a function of the transition check interval  $FcBlock$  for spawned NEBs and transition checks (+) and if only NEBs are spawned (x).

Figure 11 shows the same study using a minimal core count for the spawned work,  $TransCheckSpawnCoreCount = NEBBeadSpawnCoreCount = 1$ . The boost improvement from spawning the NEBs is still substantial; even  $NEBBeadSpawnCoreCount = 1$  is effective since it provides one core for each of the 21 NEB beads, giving a net speedup of 21/9.9, and (more importantly) it offloads the work in parallel. In contrast, while the spawned single-core transition checks do improve the boost slightly at low values for  $T_{High}$ , they actually cause a loss of boost for  $T_{High}$  above 850K, because it is quicker to do all of the transition checks in serial with 64 cores (making each force call about 9.9 times faster than a single core) than to wait for the slower single-core spawned transition check to complete once a transition does occur. This same detrimental behavior holds when both NEBs and transition checks are spawned, and thus it is better to spawn just the NEBs in this case (blue curve vs. red curve). Note that the characteristics in Figures 10 and 11 will vary a bit with the choice of  $FcBlock$  and  $FcThermalizeCheck$ , which in this case were set to 100.

**OBSERVATION 3.** *Frequent transition checks improve performance in large core-count situations, while they are detrimental in limited-core situations.*

Figure 12 shows the dependence on  $FcBlock$ , the inverse frequency of transition checks, using fixed parameters  $T_{Low} = 300K$ ,  $FcCoreCount = 1$  or 64,  $TransCheckSpawnCoreCount = 64$ , and  $NEBBeadSpawnCoreCount = 64$ . The value of  $FcThermalizeCheck$  was matched to  $FcBlock$ . We see that a serial code (red curve) gives its best performance when transition checks are infrequent—there is a maximum at  $FcBlock = 512$  (larger than 512 starts to degrade the physical predictions). In a parallel environment, at the same  $T_{High} = 1,100K$ , using  $FcCoreCount = 64$  and spawning NEBs gives a substantial speedup (green curves at  $FcBlock = 512$ ). An additional 20% can be gained by shortening the time between transition checks (solid green curve), but only if they are spawned (solid green curve compared to dashed green curve), which requires that enough cores are available. As mentioned earlier, this performance gain stems from the fact that transitions can be detected sooner, which decreases the amount of “wasted,” post-transition MD time. Looking at the  $T_{High} = 900K$  case (blue curves), the effect is essentially the same, and we note again the importance of choosing an optimal high temperature. Finally, we note that this 20% boost enhancement due to shortening the time between transition checks should increase further if the MD time is accumulated in parallel, such as by using parallel dynamics [Voter 1998], because the half-block post-transition time will be a greater fraction of the WCT between transition events. We leave investigation of this effect for future work.

## 10. SUMMARY AND FUTURE PLANS

We have introduced the concept of a parameterized application simulator, TADSim, that models the TAD AMD method as a fast-running proxy using DES. This was accomplished by abstracting out the computationally intensive force calls and replacing them by the passage of time while keeping true to the execution flow of the original TAD application. This DES simulator results in a much lower computational cost than actual TAD runs for understanding and exploring performance-related behavior. New algorithmic extensions, such as speculative spawning, can be tested before being incorporated into the actual code. Parameter scans can be used to understand and explore the limits of the current TAD implementation.

TADSim was validated against the original TAD code using an example silver material system. Parameter scans using this same physical system have produced interesting results that have allowed us to gain insight into the TAD algorithm behavior and suggest performance enhancements. We have observed a clear optimum  $T_{\text{High}}$ , given a  $T_{\text{Low}}$ , where performance is best, and we showed that  $T_{\text{Low}}$  and this optimal  $T_{\text{High}}$  come closer together with higher  $T_{\text{Low}}$ . We have evaluated the potential impact of the parallelization of TAD on performance. Adding parallelism to all force call computation is an effective way to achieve a direct speedup, assumed in the present case to give a 10-fold effect. Spawning of transition checks and NEBs further improves the performance and shifts the optimal  $T_{\text{High}}$  to a higher temperature. Increasing the frequency of transition checks can improve performance when large core counts are available. Overall, the TAD algorithm can make use of hundreds of cores. These results add value for future method development and physics research in MD.

We have seen that one of the important parameters in optimizing the TAD performance is the high temperature. A high value increases the potential speedup, whereas making it too high degrades performance by introducing excessive overhead. Shim and Amar [2011] have proposed ways to optimize the value for the high temperature on the fly in TAD simulations. Our work here complements and extends that work, in that we explore the dependence on a large number of parameters in addition to the high temperature, and we do so in the context of PDES.

Currently, TADSim supports exploration of the algorithmic choices more strongly than architectural choices. Processor speed, core counts for MD and spawning, and spawning start-up and teardown timings allow for limited architectural exploration. Future plans include coupling a more robust hardware model representing nodes, memory, and communications. Parameters will be added to support exploration of power consumption and resilience. Dynamic power usage will be tracked based on a node's characteristics and computations being performed. The value of specific resilience strategies for detection and recovery on runtime and computation quality will be explored for given hardware error rates.

A ParRep application simulator has been developed that represents the parallel-replica dynamics AMD method [Voter 1998], with publication forthcoming. Potential future directions include using TADSim for on the fly computational steering of TAD while running, exploring optimal AMD method combinations using plug-and-play component assemblies, and applying this DES approach to other application areas.

## ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their insights and helpful comments. The authors would also like to thank Kei Davis for his helpful comments. The authors would like to acknowledge the Institutional Computing Program at LANL for use of their HPC cluster resources.

## REFERENCES

- Advanced Simulation and Computing Program. 2012. National Security Applications Co-Design Project. Retrieved March 5, 2015, from <https://asc.llnl.gov/codesign/>.
- M. P. Allen and D. J. Tildesley. 1989. *Computer Simulation of Liquids*. Oxford University Press.
- P. Andelfinger and H. Hartenstein. 2013. Towards performance evaluation of conservative distributed discrete-event network simulations using second-order simulation. In *Proceedings of the 2013 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (SIGSIM-PADS'13)*. ACM, New York, NY, 221–230.
- R. Bagrodia, E. Deelman, S. Doco, and T. Phan. 1999. Performance prediction of large parallel applications using parallel simulations. In *Proceedings of the 7th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'99)*. ACM, New York, NY, 151–162.
- X.-M. Bai, A. F. Voter, R. G. Hoagland, M. Nastasi, and B. P. Uberuaga. 2010. Efficient annealing of radiation damage near grain boundaries via interstitial emission. *Science* 327, 5973, 1631–1634.
- J. Banks, J. S. Carson II, B. L. Nelson, and D. M. Nicol. 2008. *Discrete-Event System Simulation*. Prentice Hall, Englewood Cliffs, NJ.
- K. J. Barker, K. Davis, A. Hoisie, D. J. Kerbyson, M. Lang, and S. Pakin. 2009. Using performance modeling to design large-scale systems. *IEEE Computer* 42, 11, 42–49.
- K. J. Barker, S. Pakin, and D. J. Kerbyson. 2006. A performance model of the Krak hydrodynamics application. In *Proceedings of the International Conference on Parallel Processing (ICPP'06)*. IEEE, Los Alamitos, CA, 245–254.
- G. Bauer, S. Gottlieb, and T. Hoefler. 2012. Performance modeling and comparative analysis of the MILC lattice QCD application su3\_rmd. In *Proceedings of the 12th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid'12)*. 652–659.
- N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, and D. A. Wood. 2011. The gem5 simulator. *Newsletter ACM SIGARCH Computer Architecture News* 39, 2, 1–7.
- A. B. Bortz, M. H. Kalos, and J. L. Lebowitz. 1975. A new algorithm for Monte Carlo simulation of Ising spin systems. *Journal of Computational Physics* 17, 1, 10–18.
- R. N. Calheiros, M. A. S. Netto, C. A. F. De Rose, and R. Buyya. 2012. EMUSIM: An integrated emulation and simulation environment for modeling, evaluation, and validation of performance of cloud computing applications. *Software—Practice and Experience* 00, 1–18.
- T. Clark. 1985. *A Handbook of Computational Chemistry: A Practical Guide to Chemical Structure and Energy Calculations*. Wiley, New York, NY.
- M. S. Daw and M. I. Baskes. 1984. Embedded atom method: Derivation and application to impurities, surfaces, and other defects in metals. *Physical Review B* 29, 6443–6453.
- S. Eidenbenz, K. Davis, A. Voter, H. Djidjev, L. Gurvitz, C. Junghans, S. Mniszewski, D. Perez, N. Santhi, and S. Thulasidasan. 2012. Optimization principles for codesign applied to molecular dynamics: Design space exploration, performance prediction and optimization strategies. In *Proceedings of the U.S. Department of Energy Exascale Research Conference*.
- R. Ewald, J. Himmelpach, A. M. Uhrmacher, D. Chen, and G. K. Theodoropoulos. 2006. A simulation approach to facilitate parallel and distributed discrete-event simulator development. In *Proceedings of the 10th IEEE International Symposium on Distributed Simulation and Real-Time Applications (DS-RT'06)*. IEEE, Los Alamitos, CA, 209–218.
- R. M. Fujimoto. 1990. Parallel discrete event simulation. *Communications of the ACM* 33, 10, 30–53.
- D. Gianni, G. Iazeolla, and A. D'Ambrogio. 2010. A methodology to predict the performance of distributed simulations. In *Proceedings of the 2010 IEEE Workshop on Principles of Advanced and Distributed Simulation (PADS'10)*. IEEE, Los Alamitos, CA, 31–39.
- G. Hendry and A. Rodriguez. 2012. SST: A simulator for exascale co-design. In *Proceedings of the ASCR/ASC Exascale Research Conference*.
- G. Henkelman and H. Jónsson. 2001. Long time scale kinetic Monte Carlo simulations without lattice approximation and predefined event table. *Journal of Chemical Physics* 115, 9657.
- M. Jeschke, R. Ewald, and A. M. Uhrmacher. 2011. Exploring the performance of spatial stochastic simulation algorithms. *Journal of Computational Physics* 230, 7, 2562–2574.
- H. Jónsson, G. Mills, and K. W. Jacobsen. 1998. Nudged elastic band method for finding minimum energy paths of transitions. In *Classical and Quantum Dynamics in Condensed Phase Simulations*, B. J. Berne, G. Ciccotti, and D. F. Coker (Eds.). World Scientific, 385–404.
- L. Kroc, S. Eidenbenz, and V. Ramaswamy. 2007. *SimCore*. Technical Report: Los Alamos Unclassified Report 07-0590. Los Alamos National Laboratory, Los Alamos, NM.

- Q. Li, Y. Dong, D. Perez, A. Martini, and R. W. Carpick. 2011. Speed dependence of atomic stick-slip friction in optimally matched experiments and molecular dynamics simulations. *Physical Review Letters* 106, 12, 126101.
- J. Liu. 2010. Parallel discrete-event simulation. In *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley and Sons, New York, NY, 1–10.
- R. Marcelin. 1915. Contribution a l'etude de la cinetique physico-chimique. *Annales de Physique* 3, 120–231.
- J. A. McCammon and S. C. Harvey. 1988. *Dynamics of Proteins and Nucleic Acids*. Cambridge University Press.
- C. C. McGeoch. 2007. Experimental algorithmics. *Communications of the ACM* 50, 11, 27–31.
- C. C. McGeoch. 2012. *A Guide to Experimental Algorithmics*. Cambridge University Press.
- S. Mniszewski. 2010. *SimCore Tutorial*. Technical Report: Los Alamos Unclassified Report 10-04282. Los Alamos National Laboratory, Los Alamos, NM.
- Modeling & Networking Systems Research Group, Florida International University. 2012. Parallel Real-Time Immersive Network Modeling Environment (PRIME). Retrieved March 5, 2015, from <https://www.primesf.net/bin/view/Public/PRIMEProject>.
- D. Perez, B. P. Uberuaga, Y. Shim, J. G. Amar, and A. F. Voter. 2009. Accelerated molecular dynamics methods: Introduction and recent developments. *Annual Reports in Computational Chemistry* 5, 79.
- S. Plimpton. 1995. Fast parallel algorithms for short-range molecular dynamics. *Journal of Computational Physics* 117, 1, 1–19.
- D. C. Rapaport. 2004. *The Art of Molecular Dynamics Simulation*. Cambridge University Press.
- A. F. Rodrigues, K. S. Hemmert, B. W. Barrett, C. Kersey, R. Oldfield, N. Weston, R. Risen, J. Cook, P. Rosenfeld, E. Cooper-Balis, and B. Jacob. 2011. The structural simulation toolkit. *ACM SIGMETRICS Performance Evaluation Review—Special Issue on the 1st International Workshop on Performance Modeling, Benchmarking, and Simulation of High Performance Computing Systems (PMBS 10)* 38, 4, 37–42.
- N. Santhi, K. Davis, and S. Eidenbenz. 2013. *Graphical Models and a Scalable Methodology for Exascale Performance Prediction*. Technical Report: Los Alamos Unclassified Report 13-24435.
- Y. Shim and J. G. Amar. 2011. Adaptive temperature-accelerated dynamics. *Journal of Chemical Physics* 134, 054127.
- M. R. Sorensen and A. F. Voter. 2000. Temperature-accelerated dynamics for simulation of infrequent events. *Journal of Chemical Physics* 112, 21, 9599–9606.
- K. L. Spafford and J. S. Vetter. 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)*. Article No. 84.
- M. O. Steinhauser and S. Hiermaier. 2009. A review of computational methods in materials science: Examples from shock-wave and polymer physics. *International Journal of Molecular Sciences* 10, 12, 5135.
- C. Tapus and J. Hickey. 2007. A theory of nested speculative execution. In *Proceedings of the 9th International Conference on Coordination Models and Languages (Coordination'07)*. Springer-Verlag, Berlin, Heidelberg, 151–170.
- U.S. Department of Energy Office of Science. 2012. Scientific Discovery through Advanced Computing (SciDAC) Co-Design. Retrieved March 5, 2015, from <http://science.energy.gov/ascr/research/scidac/co-design/>.
- U.S. Department of Energy Office of Science. 2014. Applied Mathematics Research for Exascale Computing. Retrieved March 5, 2015, from <http://science.energy.gov/~media/ascr/pdf/research/am/docs/EMWGreport.pdf>.
- G. H. Vineyard. 1957. Frequency factors and isotope effects in solid state rate processes. *Journal of Physics and Chemistry of Solids* 3, 1, 121–127.
- A. F. Voter. 1988. Simulation of the layer-growth dynamics in silver films: Dynamics of adatom and vacancy clusters on Ag (100). In *Proceedings of SPIE 0821*. 214–226.
- A. F. Voter. 1994. The embedded atom method. *Intermetallic Compounds: Principles* 1, 77.
- A. F. Voter. 1998. Parallel replica method for dynamics of infrequent events. *Physical Review B* 57, xi224, 985–988.
- A. F. Voter. 2007. Introduction to the kinetic Monte Carlo method. In *Radiation Effects in Solids*. Springer, 1–23.
- A. F. Voter, F. Montalenti, and T. C. Germann. 2002. Extending the time scale in atomistic simulation of materials. *Annual Review of Materials Research* 32, 321.

- W. H. Wolf. 1994. Hardware-software co-design of embedded systems. *Proceedings of the IEEE* 82, 7, 967–989.
- H. Xia, H. Dail, H. Casonova, and A. Chien. 2004. The MicroGrid: Using emulation to predict application performance in diverse grid network environments. In *Proceedings of the 2nd International Workshop on Challenges of Large Applications in Distributed Environments (CLADE'04)*. IEEE, Los Alamitos, CA, 52–61.
- G. Zheng, G. Kakulapati, and L. V. Kale. 2004. BigSim: A parallel simulator for performance prediction of extremely large parallel machines. In *Proceedings of the 18th International Parallel and Distributed Processing Symposium (IPDPS'04)*. IEEE, Los Alamitos, CA.

Received November 2013; revised October 2014; accepted November 2014