

SANDIA REPORT

SAND2014-0688
Unlimited Release
Printed Jan., 2014

Supersedes SAND2013-3180
Dated April, 2013

Using Simulation to Evaluate the Performance of Resilience Strategies and Process Failures

Scott Levy, Bryan Topp, Kurt B. Ferreira, Patrick Widener, Dorian Arnold, Torsten Hoefler,

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

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

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



SAND2014-0688
Unlimited Release
Printed Jan., 2014

Supersedes SAND2013-3180
dated April, 2013

Using Simulation to Evaluate the Performance of Resilience Strategies and Process Failures

Scott Levy, Bryan Topp and Dorian Arnold
{slevy|betopp|darnold}@cs.unm.edu
Department of Computer Science
University of New Mexico
1 University of New Mexico
Albuquerque, NM 87131-0001

Kurt B. Ferreira and Patrick Widener
{kbferre|pwidene}@sandia.gov
Scalable System Software Department
Sandia National Laboratories
P. O. Box 5800
Albuquerque, NM 87185-1319

Torsten Hoefler
htor@inf.ethz.ch
Computer Science Department
ETH Zürich
Zurich, Switzerland

Abstract

Fault-tolerance has been identified as a major challenge for future extreme-scale systems. Current predictions suggest that, as systems grow in size, failures will occur more frequently. Because increases in failure frequency reduce the performance and scalability of these systems, significant effort has been devoted to developing and refining resilience mechanisms to mitigate the impact of failures. However, effective evaluation of these mechanisms has been challenging. Current systems are smaller and have significantly different architectural

features (e.g., interconnect, persistent storage) than we expect to see in next-generation systems. To overcome these challenges, we propose the use of simulation. Simulation has been shown to be an effective tool for investigating performance characteristics of applications on future systems. In this work, we: identify the set of system characteristics that are necessary for accurate performance prediction of resilience mechanisms for HPC systems and applications; demonstrate how these system characteristics can be incorporated into an existing large-scale simulator; and evaluate the predictive performance of our modified simulator. We also describe how we were able to optimize the simulator for large temporal and spatial scales—allowing the simulator to run 4x faster and use over 100x less memory.

Contents

Introduction	8
Considerations for Resilience at Scale	10
Hardware Characteristics	10
Temporal Scale	10
Spatial Scale	10
Architectural Features	10
Application Characteristics	11
Impact of Checkpoint/Restart Mechanisms	11
Impact of Failures	12
Failure Characterization	12
Restart Time	12
Recovery Model	12
LogGOPSim	13
Simulating Application Characteristics	14
Simulating Hardware Characteristics	14
Extending LogGOPSim for Large Scale Resilience Research	15
Simulating Failures and Resilience	15
Optimizing LogGOPSim for Scale	16
Evaluating our LogGOPSim Extensions	18
Correctness of the Extensions	18
Evaluating Performance Enhancements	18
Memory Usage	18

Simulation Performance	18
Validating Checkpoint Simulation With and Without Failures	19
Failure-Free Analytic Model of Coordinated Checkpointing	20
Small-scale testing	21
Validating Simulations with Node Failures	25
Related Work	26
Conclusion & Future Work	28
References	29

Introduction

Fault-tolerance has been identified as a major challenge for exascale-class systems. As systems grow in scale and complexity, failures become increasingly likely; impacting performance and scalability. Current predictions suggest that for next-generation systems the mean time between failures could fall to hours, or even minutes [14]. As failure rates increase, more time is spent preparing for and recovering from failures and less time is spent doing (useful) application work. Given these dire predictions and the dynamics of fault-tolerance techniques, significant effort has been and is being devoted to investigations aimed at improving system resilience and related mechanisms.

Effective evaluation of fault-tolerance strategies on extreme-scale systems has been difficult for several reasons. Most significantly, researchers often need to study machines that either: are larger than those that are currently available; or have hypothetical architectures or configurations. As a result, existing systems are not sufficient to evaluate the performance impact of fault tolerance techniques on next-generation extreme-scale systems. Tests performed on these systems cannot accurately account for the impact of scale and may not be able capture the impact of architectural features (e.g. interconnect technologies) whose performance varies dramatically from current systems. Second, the largest and most advanced current machines generally are not accessible to most researchers. Third, analytic techniques for predicting performance in future systems are lacking. While accurate models for coordinated checkpointing exist [7, 10], we lack analytical tools for predicting the performance impact of many other fault tolerance mechanisms (for example, message-logging [19], communication-induced checkpointing [2] and hierarchical checkpointing [27]).

The broader objective of this project is to study general fault-tolerance techniques and their impacts on application performance. However, for the work presented in this paper, we focused on *checkpoint/restart*. Checkpoint/restart (or *rollback recovery*) is the technique most commonly used on today’s systems. During normal operation, checkpoint/restart protocols [12] periodically record the state and address space of all application processes to stable storage devices. When a process fails, a new incarnation of the failed process is *recovered* from the most recent checkpoint – therefore limiting lost work. For distributed applications, coordinated checkpointing pauses all processes to record a globally consistent snapshot of the application’s state. Uncoordinated checkpointing protocols avoid synchronization overheads and I/O contention by allowing each process to checkpoint independently. Uncoordinated checkpointing protocols also avoid rolling back non-failed processes. While there have been a number of studies which show that the overheads of checkpoint/restart could be prohibitively expensive for future extreme-scale systems [15, 31, 35], there has been a great effort in the research community to optimize these rollback/recovery protocols to ensure they remain viable [8, 11, 15, 17, 19, 20, 24, 26, 29].

Researchers have shown that simulation is an effective tool for investigating the performance characteristics of applications on current and hypothetical future systems [15, 22, 30, 41]. In this paper, we focus on efficient simulation of the impact of coordinated and uncoordinated checkpoint/restart protocols on application performance. Our approach is

motivated by two observations: (1) simulation can be very computationally expensive, and simulation efficiency is maximized by considering only the features of the computing environment that are relevant to the performance impact of checkpoint/restart; and (2) the coarse-grained operation of checkpoint/restart (on the order of minutes to hours) allows us to forego the overheads and complexities of cycle-accurate simulation. Based on these observations, we hypothesize that like operating system noise [16, 21], resilience mechanisms (e.g., writing checkpoints, restarting after a failure or redoing lost work) can be modeled as CPU detours. A CPU detour is a number of CPU cycles that are used for something other than the application.

In this work, we provide a principled approach to simulating checkpoint/restart based fault-tolerance for large-scale HPC systems in a failure-prone environment. Based on this approach, we also present an efficient and accurate framework for simulating the performance impact of coordinated and uncoordinated checkpoint/restart protocols for existing and hypothetical extreme-scale systems and applications. Specific contributions of the work include:

- A survey of system, application, and resilience characteristics required for accurate and efficient simulation of extreme-scale workloads in a failure-prone environment;
- A prototype checkpoint/restart simulation framework, based on functional and performance-oriented extensions to `LogGOPSim` [22] which decrease memory consumption by over 100x and runtime by 4x;
- A validation of our hypothesis that resilience overheads can be modeled as CPU detours; and
- An evaluation of the predictive performance of our simulation approach showing an error of less than 3% against analytic models for checkpointing both with and without failures.

The organization of this paper is as follows: in the next section, we discuss the relevant system, application, failure and resilience characteristics that must be considered by our framework. Additionally, this section offers more background on checkpoint/restart protocols and shows how they factor into our considerations. The `LogGOPSim` provides an overview of the simulator that serves as the basis of our prototype. Followed by the `LogGOPSim` extension’s section, which describe the functional and performance-oriented extensions we made to `LogGOPSim` to improve its ability to simulate coordinated and uncoordinated checkpoint protocols at our desired time and space scales. Following that, we evaluate the impact of our performance optimizations and validate the accuracy of the extended simulator for checkpoint/restart. We then discuss related works, and finally, we conclude by summarizing the impact of our current contributions and plans for future work.

Considerations for Resilience at Scale

Toward the goal of efficient, large-scale simulations that allow us to evaluate resilience techniques, we must identify the relevant hardware and software characteristics that impact simulation performance. We now describe our principled approach to identifying these characteristics: in turn, we consider system features, application behavior, fault-tolerance mechanisms and then the impact of failures.

Hardware Characteristics

Our objective is to develop a simulation framework that will enable us to evaluate resilience techniques on current and future systems. The simulator must be able to accurately and efficiently model the impact of faults and fault tolerance on application performance given the: (a) temporal scale, (b) spatial scale and (c) architectural features of next-generation extreme-scale systems.

Temporal Scale

Faults and fault tolerance mechanisms typically operate at large time-scales (for example, minutes, hours or even weeks). As we stated in the introduction, projected mean-time-to-interrupt (MTTI) on the first exascale machines are on the order of hours. Additionally, many of the target applications are long running, and the behaviors of the applications as well as the systems are expected to be dynamic. As a result, simulating resilience requires a simulator that can model relatively long periods of application execution.

Spatial Scale

The largest current HPC systems are comprised of tens of thousands of nodes. If current predictions hold, the first exascale system may be over an order of magnitude larger. Our simulator must be capable of modeling the behavior of systems that are much larger than any that are currently available.

Architectural Features

The first exascale system is not projected to appear until sometime after 2020 [37]. In the intervening span of years, we expect to see advances in interconnect and persistent storage technologies. Our simulator should also allow us to evaluate the impact of these advances on resilience mechanisms.

Application Characteristics

Our simulator must be capable of accounting for the performance aspects of the applications behavior. Prior research and experience has shown that it may be sufficient to do this at the course granularity of the target application’s computation, specifically: its *communication graph*, a description of how processes communicate with each other; its *computation time*, the time between communication events; and its *dependencies*, a partial ordering of all communication and computation events. In the next section, we show how these characteristics interplay with resilience mechanisms.

Impact of Checkpoint/Restart Mechanisms

In checkpoint/restart protocols [13], the application or system saves snapshots of application state, *checkpoints*, to persistent storage. In coordinated checkpointing, all processes checkpoint at the same time (in order to mark a consistent global state), and in the event of a process failure, all processes must revert to their most recent checkpoint. While coordinated checkpoint/restart is the predominant approach, it suffers several limitations including increased overhead with system size and global process perturbations during checkpoint and recovery phases. Uncoordinated checkpoint/restart protocols, in which processes can checkpoint and recover independently, address these limitations – though they introduce new ones. In addition to these coarse protocols, many optimizations have been proposed including: diskless [32, 33, 36], hierarchical [27, 28] communication-induced [3] and incremental checkpointing [17, 18]. Despite the proliferation of resilience mechanisms, we lack effective methods for evaluating the true costs of each of these approaches on exascale systems [40].

Given the large temporal and spatial scales of the simulated systems that we wish to consider, effective simulation demands that we eliminate unnecessary detail. Existing work on modeling and simulation of coordinated checkpointing provides us with a guidepost on the required components and level of details [7, 10, 34].

In a failure-free environment, we can accurately model the impact of coordinated checkpointing by considering: the *checkpoint time*, amount of time that checkpointing activities prevent the application from executing, the *checkpoint interval*, time between consecutive checkpoints, and *work time*, the amount of time that the application would execute in the absence of checkpointing activities. Checkpoint time may need further refinement potentially including a process *coordination* phase, the *checkpoint calculation* phase during which time the checkpoint data is computed, the *checkpoint commit* time to write the checkpoint to stable storage and the *resumption* phase to continue normal application execution.

For approaches like uncoordinated checkpointing that lack explicit coordination, we also need to consider the application characteristics like communication patterns described previously. Consider a simple uncoordinated checkpointing strategy where each process generates checkpoints strictly according to local policies. Communication dependencies may cause checkpointing activities in one process may perturb the behavior and performance of

other processes. For example, if the recipient of a message is currently busy generating a checkpoint then reception of the message may be delayed until the checkpoint is complete. Further, all actions that are dependent on the reception of the message will also be delayed. Additionally, many asynchronous resilience techniques use message logging [13] to mitigate recovery costs. Accounting for this activity also requires that we incorporate information about communication patterns into our simulation.

Impact of Failures

Meaningful evaluation of resilience mechanisms necessarily includes the introduction of failures. Initially, we consider only fail-stop failures. To accurately simulate the impact of the occurrence of failures on application performance, at a minimum, we need to consider: (a) failure characterization; (b) restart time; and (c) recovery description.

Failure Characterization

To evaluate the impact of faults in the context of a resilience mechanism, we require a description of how failures occur in the simulated system. Although this could be expressed in many ways, the most common and succinct description of failure occurrences is in the form of a probability distribution.

Restart Time

When a failure occurs, some time elapses before any computation can be undertaken on the failed node. To account for this fact, we need to know the time between the occurrence of a failure and the moment when the failed node can resume computation. This includes time to restart failed nodes and processes and to read checkpoints from persistent storage, but does not include any time for recovery. For example, in the case of coordinated checkpointing, the end of the restart interval coincides with the beginning of rework (i.e., redoing work lost due to the failure).

Recovery Model

When the failed node has restarted and is able to resume computation, there is typically some amount of work that needs to be redone before the system can again make meaningful forward progress. For example, in coordinated checkpointing, all of the computation between the last valid checkpoint and the occurrence of the failure needs to be redone. Typically, each resilience mechanisms presents a different method for recovering from a failure. Therefore, to accurately account for the cost of recovering from a failure, we need a model for each

Required to Model	Parameter Name	Parameter Description
All Checkpointing	COORDINATION TIME	time for processes to coordinate the taking of a checkpoint
	CHECKPOINT COMPUTATION	time to compute a checkpoint
	CHECKPOINT COMMIT TIME	time to write a checkpoint to stable storage
	CHECKPOINT INTERVAL WORK TIME	time between consecutive checkpoints time-to-solution without failures or resilience mechanisms
Uncoordinated Checkpointing	COMMUNICATION GRAPH COMPUTATION EVENTS	details of inter-process communication failure-free computation pattern of the application
	DEPENDENCIES	partial ordering of communication and computation events
Failure Occurrences	FAILURE CHARACTERIZATION	rate and distribution of failures
	RESTART TIME	time to read a checkpoint from stable storage after a failure
	RECOVERY MODEL	a model of the time required before forward progress can resume

Table 1. Summary of the parameters needed for accurate simulation of HPC applications in a failure-prone system.

resilience mechanism that allows us to determine the amount of time that will elapse before the application resumes forward progress.

LogGOPSim

In this section, we describe `LogGOPSim` [22, 38], the simulator we extend to meet the requirements prescribed by the considerations in the previous section. We choose `LogGOPSim` because it is shown to be accurate, freely available and fast enough to support large-scale simulations while already capturing many of the application and hardware characteristics we require (as we discuss). As described in Section , functionally, we simply needed to extend it to account for checkpoint/restart and failure recovery.

Simulating Application Characteristics

LogGOPSim is an application simulator based on the popular LogP model [9]. LogP and its variants have a long history of accurately predicting the performance of large-scale parallel applications and algorithms. The simulation framework consists of three major components: a trace collector (`liballprof`), a schedule generator (`SchedGen`), and an optimized discrete-event simulator (`LogGOPSim`).

The trace collector records the actual MPI communication of the target application. The schedule generator uses the MPI traces to generate a schedule that captures the required characteristics of control- and dataflow of the application while preserving the happens-before relationship of events within the application. The discrete-event simulator reads the generated schedule, performs a full LogGOPS simulation and reports the completion time of each process.

This validated simulation framework was developed to simulate applications at scale, and has the ability to extrapolate from traces collected on smaller scale systems. This allows for the simulation of platforms larger than those currently in existence while keeping the same communication characteristics (equivalent to weak-scaling an application). Although the extrapolated trace may not precisely represent the communication pattern on the larger system, the impact of this inaccuracy has been shown to be small [22] if extrapolation factors are bounded. This framework has been used to evaluate the performance of collective communications [23] and the impact of OS noise [21] on large-scale applications. A detailed study of the simulation framework and its functionality is presented in [22].

Simulating Hardware Characteristics

Because LogGOPSim was initially developed to model application performance in large-scale systems [21], it allows us to model systems with the characteristics described in the preceding section. First, it provides the simulation scale necessary for evaluating checkpointing techniques. For a single collective operation, LogGOPSim can simulate up to 10,000,000 processes. For more general workloads, it is capable of simulating more than 64,000 processes.

Second, with some minor modifications, LogGOPSim is also capable of simulating the necessary temporal scale. The initial implementation of LogGOPSim was intended for comparatively short simulations. As a result, the temporal scope of the simulations that can be executed by the unmodified simulator is significantly limited by the size of the simulating system’s memory. To achieve the temporal scale that we needed with reasonable quantities of system memory, we made some simple modifications to LogGOPSim. These modifications are discussed more fully in a subsequent section.

Third, LogGOPSim also allows us to model the impact of emerging interconnect technologies. Working within the LogGOPS model, we can simulate the impact of many changes in network behavior on resilience techniques by modifying the model’s parameters. In addition,

as we discuss more fully below, our model of resilience mechanisms allows us to evaluate how improvements to persistent storage systems (e.g., the widespread availability of local SSDs) will affect the performance of resilience mechanisms.

Extending LogGOPSim for Large Scale Resilience Research

Simulating Failures and Resilience

The key insight that allows us to use LogGOPSim is that resilience mechanisms (e.g., writing checkpoints, restarting after a failure, redoing lost work) can be modeled as CPU *detours*. A CPU detour is a number of CPU cycles that are used for something other than the application, similar to OS noise [16, 21]. One key difference between OS noise and these resilience detours is that resilience “noise” events may need to be replayed synchronously with the application communication/computation pattern rather than asynchronously as is typical of OS noise.

We model resilience in LogGOPSim using a new library, `libsolipsis`, that generates CPU detours based on a specified resilience mechanism and the application’s communication pattern. Similar to `liballprof`, the library links to the application using the MPI profiling interface, intercepting all MPI calls. The output of this library is a per-process detour file that can be provided as input to LogGOPSim. The detour file contains the timestamp and the duration of each of the resilience mechanism detours. The duration of detours, T_{detour} , that represent checkpoints are computed using the following expression.

$$T_{detour} = T_{coord} + T_{ckpt} + T_{commit}$$

where

T_{coord} = time to coordinate the taking of a checkpoint

T_{ckpt} = time to compute a checkpoint

T_{commit} = time to commit the checkpoint to stable storage

We also generate detours to represent the impact of node failure and pessimistic message logging. In the case of failure, the duration of the detour includes both the restart and rework time on the failed node; `libsolipsis` computes the rework time by calculating the amount of simulated time that has elapsed since the previous checkpoint. For pessimistic message logging, `libsolipsis` calculates the time required to write the message to the log given a bandwidth to stable storage.

For the purposes of this work, we focus on the libraries’ ability to emulate performance of two popular resilience mechanisms: coordinated checkpointing and asynchronous checkpointing with message logging [12].

We focus on these two methods because coordinated checkpoint/restart is currently the most popular approach and asynchronous checkpointing has been proposed as a low-overhead checkpoint option for future extreme-scale systems.

For asynchronous checkpointing with message logging, our library generates detour files that contain the timestamp and the duration of the local checkpoints. Because no coordination is required, $T_{coord} = 0$. Also, for simplicity, we currently assume that $T_{ckpt} = 0$. For pessimistic message logging [12], we modify the CPU overhead parameter (o in the LogGOPS model) for send operations (o_s) to account for the log write to stable storage. The LogGOPSim simulator uses a single detour file to simulate the local checkpoints in the system; to model the asynchronous nature of these checkpoints, each node starts at a different location in the file.

For coordinated checkpoint/restart, the library generates a detour file that contains the timestamp and the duration of each checkpoint taken by the application. For this work, we have assumed bulk-synchronous parallel (BSP) applications. Because applications of this type are largely self-synchronizing, we set $T_{coord} = 0$. And again, for simplicity, we are currently assuming that $T_{ckpt} = 0$. When the simulation is run, we use the “`--noise-cosched`” option of the LogGOPSim simulator. This option ensures that the detour file is co-scheduled on all processors, thereby simulating coordinated checkpoint/restart. We also force each process to start at the beginning of the detour file to ensure proper timing of checkpoints.

To simulate failure, the library generates failure times for each node from a random distribution based on a per-node mean time between failure (MTBF). When a failure is generated, the library adds a detour event that includes the the time required to restart from the last checkpoint and the time required for rework (i.e., the time since the last checkpoint). The LogGOPSim simulator will ensure that all communication in the trace file that depends on the failed node will be delayed until the node has “recovered”.

Optimizing LogGOPSim for Scale

To simulate periods of execution long enough to be meaningful for fault tolerance (i.e., application wallclock times long enough that application failures are expected) while keeping traces manageable, we extended LogGOPSim to support automatic execution trace extrapolation. Because LogGOPSim was originally designed to simulate single collective operations and short application traces it assumed a comparatively small input dataset. In our use cases, the extent to which the existing LogGOPSim could scale up the length of simulated execution and the number of simulated nodes was severely limited by the amount of available memory.

LogGOPSim, as originally published, requires a pre-processing step which performs the extrapolation to generate communication data for all simulated nodes. The simulator binary then attempts to map this file into virtual memory and use it directly as input data about simulated events. The size of this file is proportional to both the length of the simulated execution and the number of simulated nodes. As a result, simulating long running applica-

tions or large-scale systems requires very large data sets. Additionally, when collecting data at varying scales, a user would be required to re-run the entire toolchain from the trace data to the simulator with different parameters.

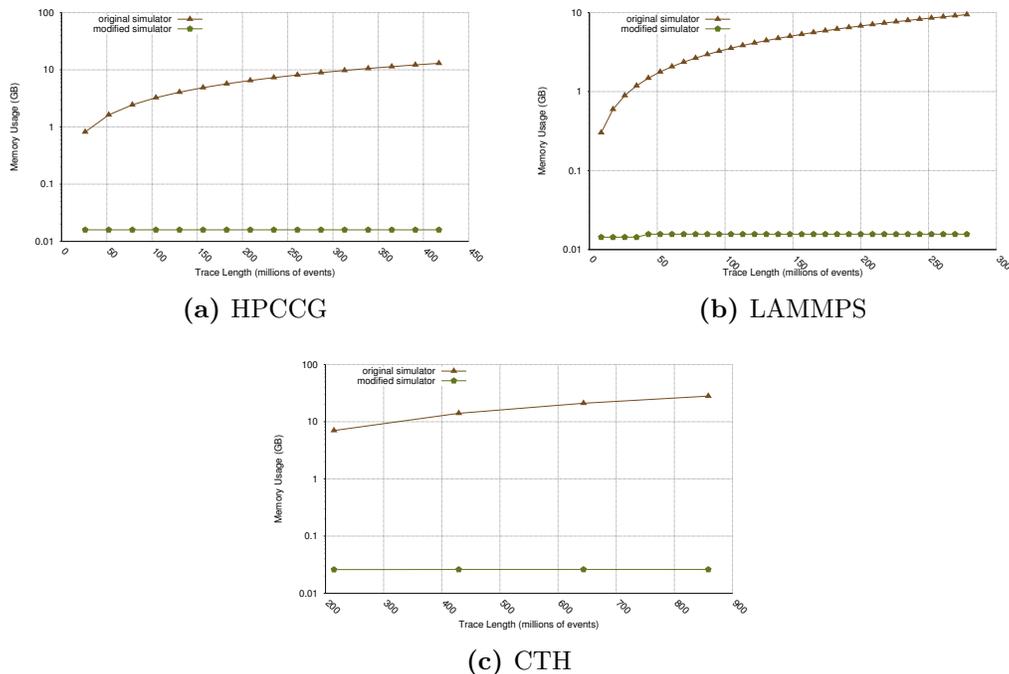


Figure 1. Comparison of the memory consumption required to simulate a system running one of three applications using the original `LogGOPSim` simulator and our modified version as a function of input trace length. Our windowing protocol decouples memory usage from trace length. As a result, with a fixed memory budget, our modifications allows us to simulate much longer periods of application execution than was possible with the original simulator.

We re-wrote the input handling portion of `LogGOPSim` to include two critical changes. First, the modified simulator performs extrapolation in main memory as needed, rather than as a pre-processing step on disk. The traces generated from profiling an MPI application are used directly as input, and are proportional in size to the original profiled node count, rather than the extrapolated node count. Second, the simulator works on a small sliding window of input data, rather than mapping it in all at once. The code loads and extrapolates data from the traces at fine granularity, loading only a small portion of the trace file at a time. Because of these changes, the simulator’s memory usage, shown in Figure 1, remains constant independent of input trace size. In other words, the same amount of memory would be required to simulate a minute, hour, day, week or month of application execution time!

In theory, an operating system should be able to perform this type of efficient memory

allocation when using a system call such as `mmap`. However, on the Linux 2.6.32 systems that we used, windowing the input data at the application level allowed for much greater scales before the system started to thrash virtual memory.

Evaluating our LogGOPSim Extensions

Correctness of the Extensions

We have verified with a careful comparison of the sequences of simulation events generated by each of the two simulators that our modified LogGOPSim produces exactly the same sequence and timing of simulated events as the original validated LogGOPSim tools. Moreover, the two simulators produce identical simulated runtimes.

Evaluating Performance Enhancements

In this section, we evaluate the performance impact of our modifications to LogGOPSim. We consider two important metrics to evaluate our changes, maximum memory requirement and simulation performance in events/second, and then finally examine overall wallclock time for simulating the same problem

Memory Usage

With our changes, the amount of disk space needed is no longer proportional to the node extrapolation factor, and the amount of RAM needed is no longer proportional to the length of the trace data. This enables simulation of long executions of many nodes: for the traces used here, memory usage decreased dramatically as shown in Figure 2. Memory usage dropped by a factor of 20 for HPCCG, 60 for LAMMPS, and 900 for CTH, with the magnitude of the drop related to an applications communication pattern and the greatest distance between the initiation of a non-blocking operation and waiting for its completion. This increase of available memory allowed us to simulate over 12 minutes of HPCCG at 256K nodes and over 7 minutes of LAMMPS at 256K nodes in a short amount of time, as shown in Figures 2(a) and 2(b), respectively.

Simulation Performance

Figure 3 and Figure 4 show the increase in performance for our simulation framework. We show this increase both in terms of event per second of the simulator (Figure 3) and the wall clock time to perform the simulation (Figure 4). We see from these figures, a factor of 2.5 to 4X increase in performance from our modifications. We believe that the substantial

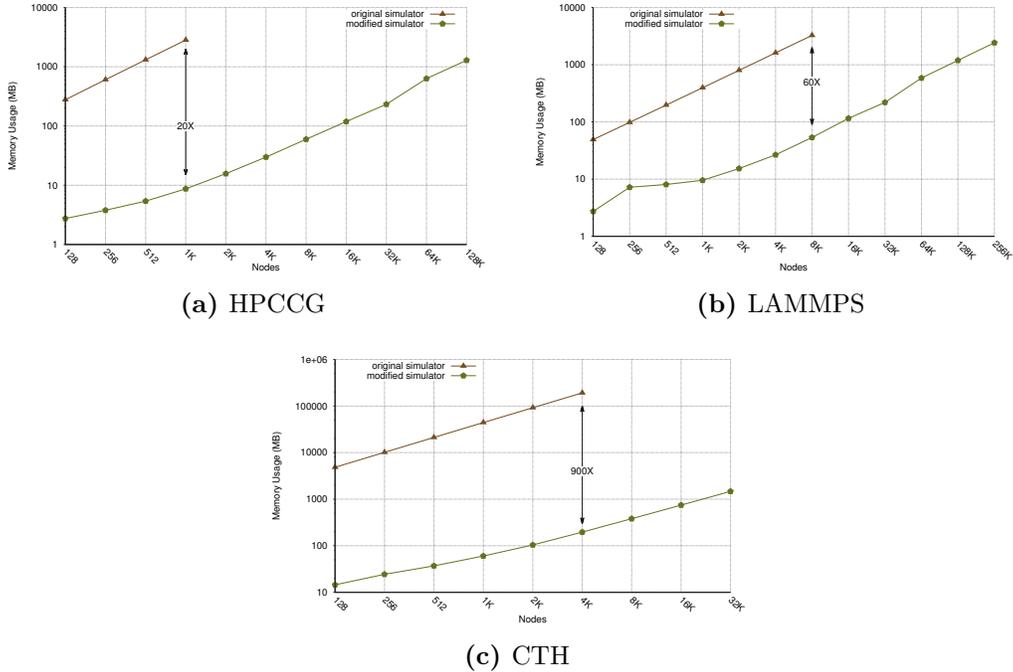


Figure 2. Comparison of the memory consumption required to simulate a system running one of three applications using the original `LogGOPS` simulator and our modified version. With a fixed memory budget, our modifications allows us to simulate systems that are significantly larger than could be simulated with the original simulator. The memory consumption decrease varies by communication pattern and varies from 20X for HPCCG to 900X for CTH.

performance benefits stem from the smaller cache footprint of our implementation. We note that simulation performance decreases slightly as the number of simulated nodes increases. We are working on characterizing this decrease. However, we conclude that the achieved performance is sufficient for our purposes.

Validating Checkpoint Simulation With and Without Failures

In this section, we present the data we collected to validate our simulator. We use both analytic models and small-scale testing to ensure that our simulator accurately models the impact of resilience mechanisms in failure-free and failure-prone environments.

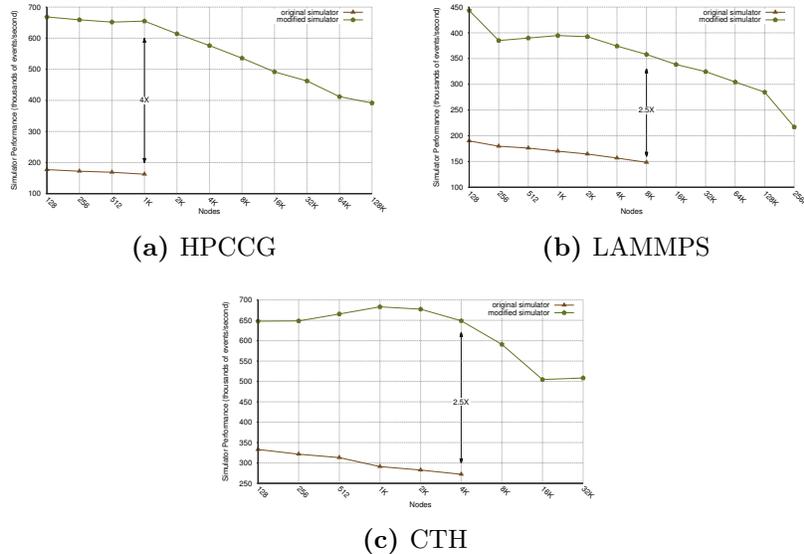


Figure 3. Comparison of simulator performance, measured in events/second, when simulating a system running one of three applications using the original LogGOPSim simulator and using our improved version. Due to the memory requirements of the original simulator, we were unable to obtain results for simulations of large-scale systems using the original simulator. The simulation performance increase varies by application communication pattern and varies from 2.5X for LAMMPS and CTH to 4X for HPCCG.

Failure-Free Analytic Model of Coordinated Checkpointing

We begin with a simple analytic model for coordinated checkpointing. Equation 1 models application performance in terms of its wall clock time-to-solution, T_w , in a failure-free environment.

$$T_w = T_s + \frac{T_s}{\tau} \times \delta \quad (1)$$

where T_w is the wall clock time, T_s is the solve time of the application without any resilience mechanism, τ is the checkpoint interval [10], and δ is the checkpoint commit time (time to write one checkpoint). For coordinated checkpointing to shared stable storage, we can express the checkpoint commit time as:

$$\delta = \frac{N * ||c_{avg}||}{\beta} \quad (2)$$

where N is the number of nodes, $||c_{avg}||$ is the average checkpoint size per node, and β is the aggregate write bandwidth to stable storage.

In Figures 5(a) and 5(b), we compare the output of this model to the output of our

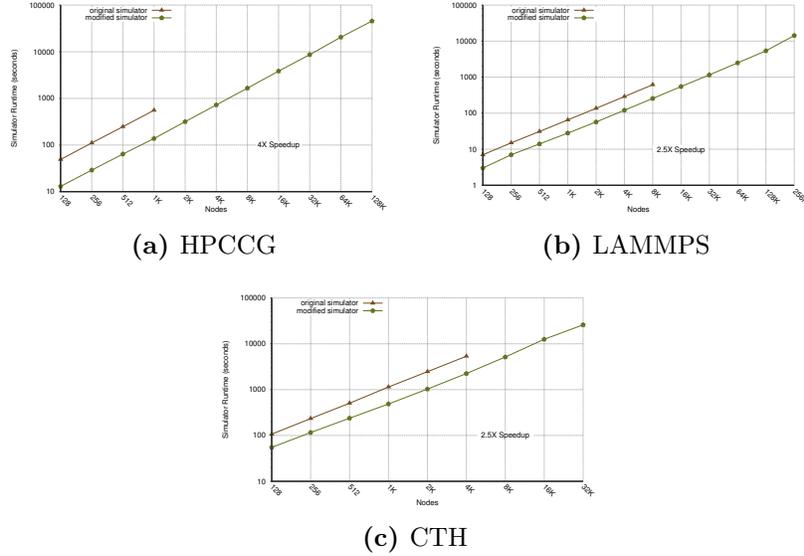
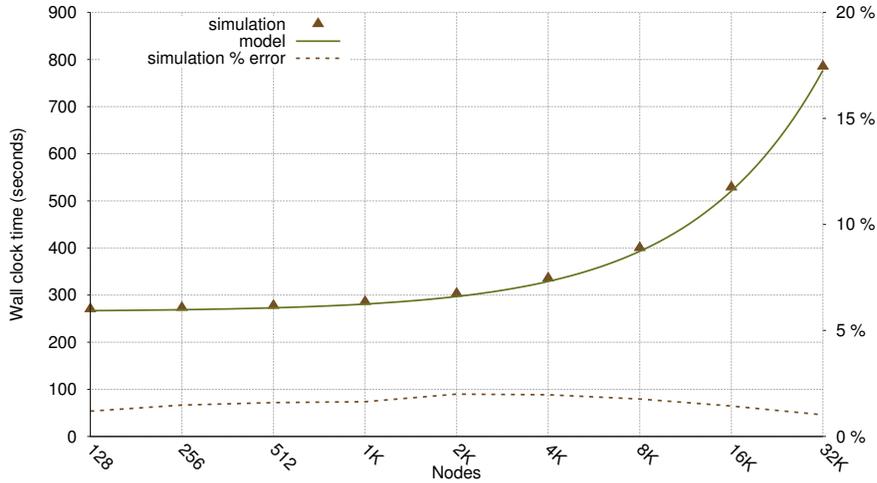


Figure 4. Comparison of simulator runtime, measured in seconds (lower is better), for the original LogGOPSim simulating systems running three different applications using simulator and using our improved version. Due to the memory requirements of the original simulator, we were unable to obtain results for simulations using the original simulator for large-scale systems. Similar to Figure 3, the simulation wall-clock speedup varies by application; from 2.5X for LAMMPS and CTH to 4X for HPCCG.

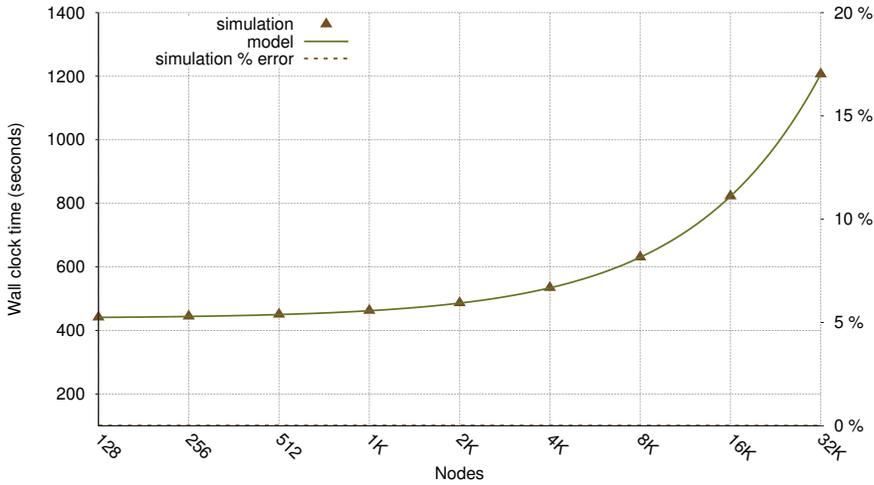
simulator. The times-to-solution for CTH predicted by the simulator are very accurate, about 3% greater than the model’s predictions. More importantly, the simulator closely matches scaling trends predicted by the model. Moreover, the simulated times-to-solution for LAMMPS are within 1% of the analytic model. On the whole, these data suggest that the simulator is accurately modeling how the impact of resilience mechanisms scales with system size.

Small-scale testing

To further validate our simulator, we compared it against the results of small-scale tests on real hardware. The simulator provides us with fine-grained control over the checkpoint interval and duration. To mimic this degree of control on real hardware, we constructed an MPI profiling library, `libchkpt`. This library, based on the the `libhashchkpt` incremental checkpointing library [17], also has the ability to take both full coordinated and uncoordinated checkpointing techniques, in addition to its incremental coordinated techniques. The full coordinated checkpointing functionality ensures all checkpoints are taken simulta-



(a) CTH

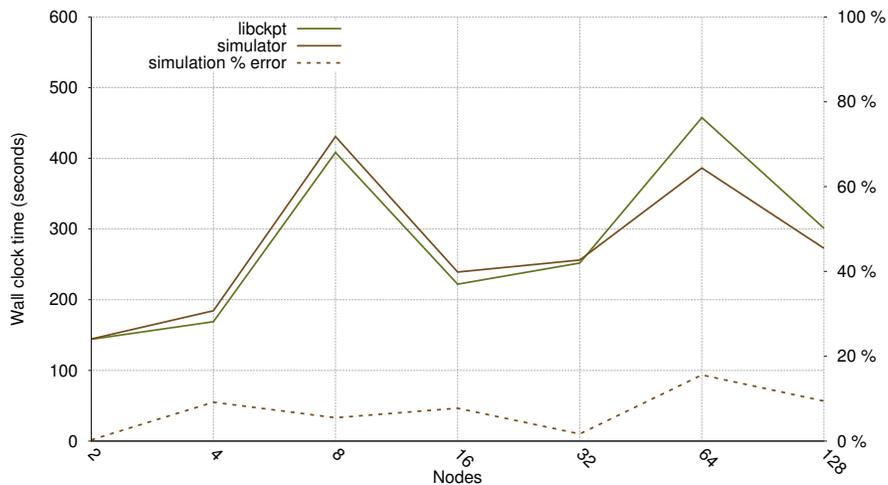


(b) LAMMPS

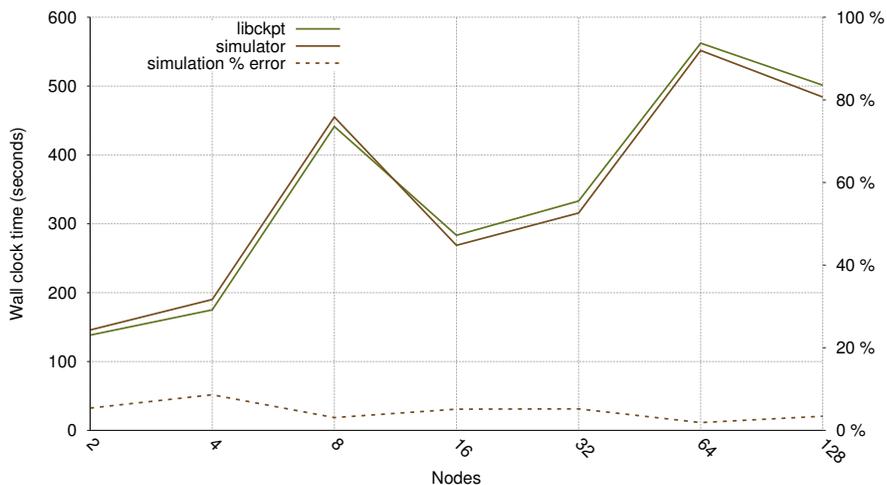
Figure 5. Validation of the simulator against the simple analytic model described in Equation 1 for coordinated checkpointing to stable storage in a failure-free environment for CTH and LAMMPS. The model and the simulator use identical values for the T_s (for each application), τ , and δ . The simulation error is less than 3% for CTH and less than 1% for LAMMPS across the tested node count range.

neously on each node, while the uncoordinated approach takes checkpoints independently. While taking checkpoints, the CPU is taken from the application until the checkpoint commit time has completed.

For our purposes here, we focus on validating the failure-free case. Figure 6 and Figure 7

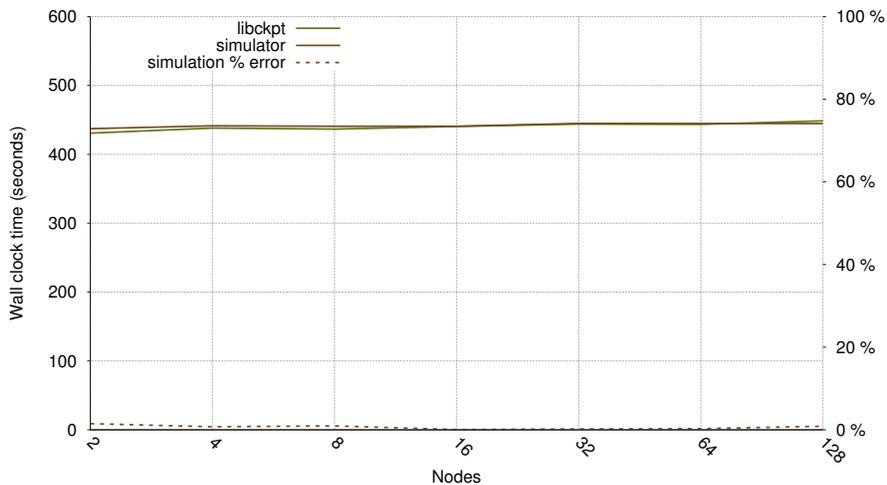


(a) Coordinated Checkpointing

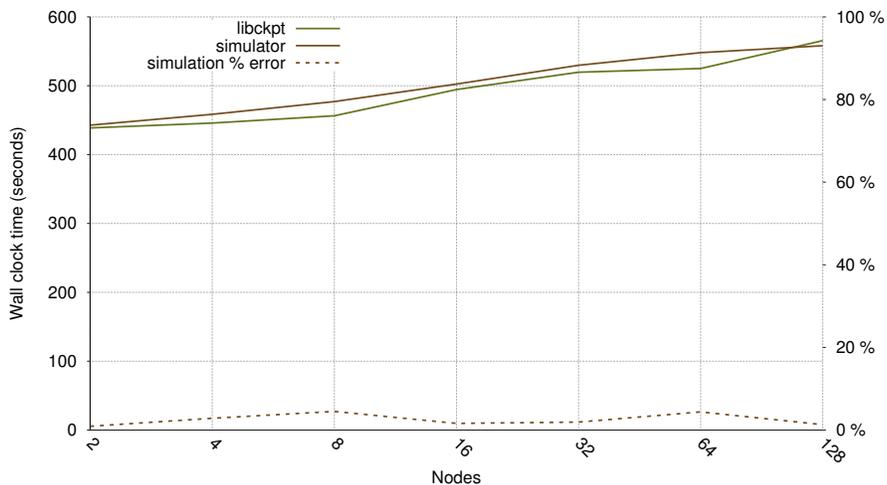


(b) Uncoordinated Checkpointing

Figure 6. Performance of LogGOPSim simulation against a coordinated and uncoordinated checkpointing library for CTH. The simulator and `libckpt` use identical values for T_w (failure free performance), τ (checkpoint interval), and δ (checkpoint commit time). The simulation error in this figure is less than 20%, with this differences attributed to platform features not being simulated. For example, interference from the OS is not being generated in this case to simplify analysis. This OS interfere has been shown to greatly influence impact CTH performance [16].



(a) Coordinated Checkpointing



(b) Uncoordinated Checkpointing

Figure 7. Performance of LogGOPSim simulation against an coordinated and uncoordinated checkpointing library for LAMMPS. The simulator and `libckpt` use identical values for T_w (failure free performance), τ (checkpoint interval), and δ (checkpoint commit time). The simulation error in this figure is shown to be less than 5% in the range tested.

show the results of this validation. These figures compare the total wall clock time simulated by `LogGOPSim` and measured with `libchkpt` running on our test platform. For reference, each figure also includes the total wall clock time in the absence of any failures. Note the performance of CTH in Figure 6 exhibits a distinct sawtooth pattern. This pattern is an artifact of how CTH scales the computation as nodes counts increase. The simulator accurately predicts this complex sawtooth pattern. We also see in this figure the simulators error in prediction. We also note that the predictive performance of the simulator is less accurate for CTH in comparison to LAMMPS, with the error in time to solution bounded by 20%. This is due to the simulator not accounting for OS noise on the node and limitations of the network model used. In our testing, OS interference is not being generated to simplify analysis, though the simulator allows for such accounting. This OS interfere has been shown to greatly influence impact CTH performance [16]. Also, though the `LogGOPSim` simulator is capable of sophisticated network models, in this work we use a simple network model which does not account for network contention. As CTH does a fair amount of bulk data transfer, network contention can be an issue.

Overall, these figures show that `LogGOPSim` closely tracks the results measured with `libchkpt`. For all the configurations that we examined, the absolute wall clock time simulated by `LogGOPSim` is within 20% of the measured values. More importantly, `LogGOPSim` closely mimics the trends we observe with `libchkpt` even as performance deviates from performance on actual hardware.

Validating Simulations with Node Failures

Finally, we validate the performance of our simulator in an environment that includes process failures. Analytic models for checkpoint/restart (CR) enable the prediction of execution times in the case of failures. For a serial computation, one can assume that the time to execute an application using CR can be modeled by:

$$T = T_{\text{app}} + T_{\text{c}} + T_{\text{rework}} + T_{\text{r}} \quad (3)$$

Here T_{app} is the time of the application execution without any resiliency mechanisms, T_{c} and T_{r} are the cumulative costs of all checkpoints and restarts, respectively, and T_{rework} is the cumulative cost of all work that was performed after the last checkpoint before each failure. Let M , the Mean Time To Interrupt (MTTI), be the mean of a Poisson distributed random variable that models component failures in a system and δ the time to perform a checkpoint. Daly [10] showed that the serial application execution time can be estimated as

$$T(\tau) = T_{\text{app}} + (k - 1)\delta + k \left(\frac{\tau + \delta}{2} + R \right) \left(\frac{\tau + \delta}{M} \right) \quad (4)$$

for a recovery time R and a checkpoint interval τ (assuming $T_{\text{app}} = k\tau, k \in \mathbb{N}$).

Assuming $\delta < 2M$, this formula can be used to derive the optimal checkpointing interval, $\hat{\tau}$ [10].

$$\hat{\tau} = \sqrt{2\delta M} \left[1 + \frac{1}{3} \sqrt{\frac{\delta}{2M}} + \frac{1}{9} \left(\frac{\delta}{2M} \right) \right] \quad (5)$$

If we assume that all component failures are identical, then $\hat{\tau}$ is also the optimal checkpointing interval for each node of a parallel computer. We next discuss how to model coordinated and uncoordinated checkpoint protocols.

Coordinated CR can be modeled using an extended BSP model with additional checkpoint supersteps that are performed after a communication superstep ends (when the network is quiet). The synchronous nature of all supersteps (phases) allows us to model the runtime of an application as $T(\hat{\tau})$ using Equation (4).

As described previously, our simulation is able to simulate scenarios that include failures and failure recovery. As with failure-free simulations, we validated the simulator’s ability to model the impact of failures by comparing against the same validated analytical model in Equation 4.

Figure 8 plots the results for simulated and modeled ten hour executions of LAMMPS using a ten minute node restart time, a one second checkpoint commit time and five year node MTBF. We use the optimal checkpoint frequency, 68 seconds, computed using the method in Equation 5. These results show that our simulator produces highly accurate output compared to the predictions of the analytical model.

Related Work

Although fault tolerance for HPC has been a very active area of research, few tools exist that allow us to project behavior beyond small-scale systems. As we discussed above, simulating fault tolerance techniques requires an appropriate level of detail about the communication of the target application. Without an accurate representation of application communication, we cannot accurately simulate some fault tolerance techniques (e.g., asynchronous checkpointing). Too much detail unnecessarily reduces simulator performance. The application simulators for fault tolerance that do exist tend to fall to either extreme; either they are not communication-accurate or they simulate communication in much greater detail than believed necessary.

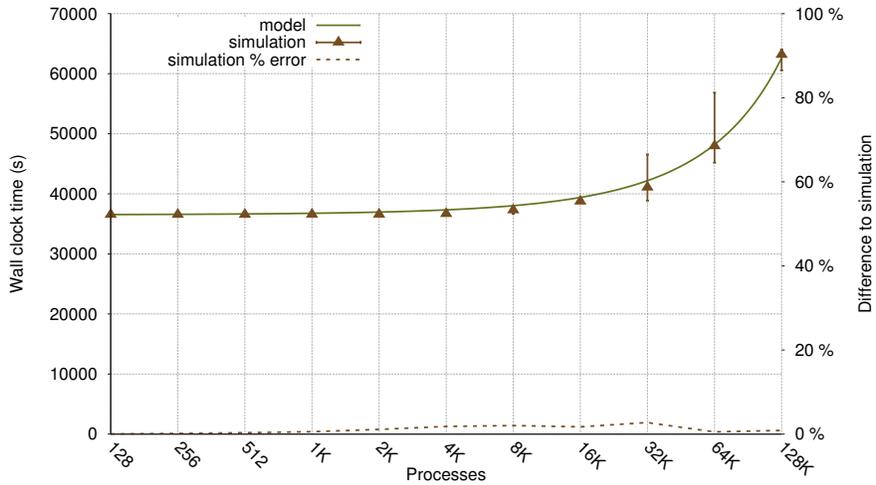


Figure 8. A comparison between efficiency of LAMMPS predicted by the simulator and by a validated analytical model (Equation 4) in the presence of failures. The simulator data for each system size represents the range of values produced over 16 runs of the simulator. This simulation models a 10.1 hour run of LAMMPS using a node restart time of 10 minutes, a checkpoint commit time of 1 second, and a node MTBF of 5 years. Checkpoints are taken every 68 seconds based on the optimal checkpoint interval described in Equation 5.

In [34], Riesen et al. present a simulator that can model the impact of node failure on application performance in the context of traditional coordinated checkpoint/restart. This simulator can also account for process replication. Tikotekar et al. take a similar approach in [39]. They present a simulator that models coordinated checkpointing and can also simulate fault prediction and process migration. While these tools have been shown to be effective for their stated purposes, they are not communication-accurate. As a result, they are unable to account for fault tolerance techniques whose performance may be influenced by communication patterns.

At the other extreme is xSim [5]. xSim builds on the MPI profiling interface and interposes itself between the application and the MPI library. As a result, the simulator is able to run unmodified HPC applications. Scaling is achieved by oversubscribing the nodes of the system used for validation. While this provides us with a tremendous amount of detail about the performance of the application, it imposes a significant cost. Due to limits on the degree of oversubscription, large-scale systems are required to simulate systems that approach extreme-scale. Moreover, as the size of the simulated system grows and the degree of oversubscription therefore increases, the time required to simulate the system grows dramatically. Lastly, this oversubscription could place significant limits on the size of the problem that can be solved as the memory for each simulated node must exist in the memory of one physical node. In contrast, our approach allows us to simulate fault tolerance mechanisms for systems comprised of tens or hundreds of thousands of nodes on very modest hardware (e.g., a single node). In some cases, this simulation completes in less time than it would take to run the application itself, but with the less detail of the computation.

Boteanu et al. present a fault tolerance extension to an existing simulator in [6]. However, they target a datacenter environment where each job is a discrete unit that is assigned to a single processing element. As a result, their simulator does not map well to HPC workloads.

Finally, SST/macro [1,25] is a coarse-grained, lightweight simulator designed to simulate the performance of existing and future large-scale systems. By collecting traces of application execution, SST/macro is able to simulate the application’s computation and computation patterns at scales and on hardware that does not yet exist. However, SST/macro does not currently account for the impact of CPU detours (OS jitter). As a result, because the foundation of our approach is based on the observation that resilience can be modeled as CPU detours, we concluded that SST/macro was not a suitable starting point for our investigation.

Conclusion & Future Work

We presented in this work, a new and promising approach to simulation at scale of fault-tolerance mechanisms based on the checkpoint/restart model. We identified a set of platform, application, and resilience characteristics required for accurate and efficient simulation; described a prototype framework based on extensions to a validated and freely-available appli-

cation simulator implementing the LogP model; shown how resilience processing overheads can be effectively modeled as CPU detours; and demonstrated empirically that our approach accurately predicts, with an error of less than 3%, the impact of resilience mechanisms. Our modifications to the LogGOPSim simulator greatly decreased its memory consumption by a factor of 100 or greater and runtime by a factor of 4. This performance increase allows us to evaluate potential resilience solutions at meaningful application and temporal scales, while also enabling the modeling of future interconnection and storage technologies.

The design space for evaluating resilience methods in large-scale HPC applications is young and still evolving. While our simulation framework has expanded that space in new and useful ways, several areas for future work remain. Among these, we intend to extend the framework to provide injection of additional failure types beyond fail-stop for large-scale simulation; e.g. corruption of application memory or network traffic. We understand that this may mean re-evaluating the granularity of our simulation to ensure proper and effective simulation. We are also investigating mechanisms to integrate both coarse- and fine-grained simulation for failures. This will, for example, allow us to use coarse-grained simulation in areas where failures occur, and fine-grained simulation when failures or other interesting events do occur. We also plan to address support for additional resilience mechanisms such as hierarchical checkpointing, replication-based approaches, process migration and cloning, as well as integration with ongoing standards efforts like the current fault tolerance proposal put forth in the MPI forum [4]. Finally, we plan to further investigate performance limitations of our current simulation framework, including analyzing the benefit of parallelization.

References

- [1] Sst: The structural simulation toolkit. http://sst.sandia.gov/about_sstmacro.html, 2011.
- [2] L. Alvisi, E. Elnozahy, S. Rao, S.A. Husain, and A. de Mel. An analysis of communication induced checkpointing. In *Fault-Tolerant Computing, 1999. Digest of Papers. Twenty-Ninth Annual International Symposium on*, pages 242–249, 1999.
- [3] Lorenzo Alvisi, Elmootazbellah Elnozahy, Sriram Rao, Syed Amir Husain, and Asanka De Mel. An analysis of communication induced checkpointing. In *Fault-Tolerant Computing, 1999. Digest of Papers. Twenty-Ninth Annual International Symposium on*, pages 242–249. IEEE, 1999.
- [4] Wesley Bland, Aurelien Bouteiller, Thomas Herault, Joshua Hursey, George Bosilca, and Jack J. Dongarra. An evaluation of user-level failure mitigation support in mpi. In Jesper Larsson Träff, Siegfried Benkner, and Jack J. Dongarra, editors, *Recent Advances in the Message Passing Interface*, volume 7490 of *Lecture Notes in Computer Science*, pages 193–203. Springer Berlin Heidelberg, 2012.
- [5] S Bohm and Christian Engelmann. xSim: The extreme-scale simulator. In *High Performance Computing and Simulation (HPCS), 2011 International Conference on*, pages 280–286. IEEE, 2011.
- [6] Adrian Boteanu, Ciprian Dobre, Florin Pop, and Valentin Cristea. Simulator for fault tolerance in large scale distributed systems. In *Intelligent Computer Communication and Processing (ICCP), 2010 IEEE International Conference on*, pages 443–450. IEEE, 2010.
- [7] Mohamed-Slim Bouguerra, Thierry Gautier, Denis Trystram, and Jean-Marc Vincent. A flexible checkpoint/restart model in distributed systems. In *Parallel Processing and Applied Mathematics*, pages 206–215. Springer, 2010.
- [8] Greg Bronevetsky, Daniel Marques, Keshav Pingali, Sally McKee, and Radu Rugina. Compiler-enhanced incremental checkpointing for openmp applications. In *IEEE International Symposium on Parallel&Distributed Processing*, pages 1–12, 2009.
- [9] David Culler, Richard Karp, David Patterson, Abhijit Sahay, Klaus Erik Schauer, Eunice Santos, Ramesh Subramonian, and Thorsten von Eicken. Logp: towards a realistic model of parallel computation. *SIGPLAN Not.*, 28(7):1–12, July 1993.
- [10] J. T. Daly. A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Gener. Comput. Syst.*, 22(3):303–312, 2006.

- [11] Xiangyu Dong, Naveen Muralimanohar, Norm Jouppi, Richard Kaufmann, and Yuan Xie. Leveraging 3d pcam technologies to reduce checkpoint overhead for future exascale systems. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, SC '09*, pages 57:1–57:12, New York, NY, USA, 2009. ACM.
- [12] E. N. (Mootaz) Elnozahy, Lorenzo Alvisi, Yi-Min Wang, and David B. Johnson. A survey of rollback-recovery protocols in message-passing systems. *ACM Comput. Surv.*, 34(3):375–408, 2002.
- [13] Elmootazbellah N. Elnozahy, Lorenzo Alvisi, Yi-Min Wang, and David B. Johnson. A survey of rollback-recovery protocols in message-passing systems. *ACM Computing Surveys*, 34(3):375–408, 2002.
- [14] Keren Bergman et al. Exascale computing study: Technology challenges in achieving exascale systems. [http://www.science.energy.gov/ascr/Research/CS/DARPAexascale-hardware\(2008\).pdf](http://www.science.energy.gov/ascr/Research/CS/DARPAexascale-hardware(2008).pdf), September 2008.
- [15] Kurt Ferreira, Rolf Riesen, Patrick Bridges, Dorian Arnold, Jon Stearley, James H. Laros III, Ron Oldfield, Kevin Pedretti, and Ron Brightwell. Evaluating the viability of process replication reliability for exascale systems. In Scott Lathrop, Jim Costa, and William Kramer, editors, *SC*. ACM, November 2011.
- [16] Kurt B Ferreira, Patrick Bridges, and Ron Brightwell. Characterizing application sensitivity to os interference using kernel-level noise injection. In *Proceedings of the 2008 ACM/IEEE conference on Supercomputing*, page 19. IEEE Press, 2008.
- [17] Kurt B. Ferreira, Rolf Riesen, Ron Brightwell, Patrick G. Bridges, and Dorian Arnold. Libhashckpt: Hash-based incremental checkpointing using GPUs. In *Proceedings of the 18th EuroMPI Conference*, Santorini, Greece, September 2011 [to appear].
- [18] Roberto Gioiosa, Jose Carlos Sancho, Song Jiang, Fabrizio Petrini, and Kei Davis. Transparent, incremental checkpointing at kernel level: a foundation for fault tolerance for parallel computers. In *Proceedings of the 2005 ACM/IEEE conference on Supercomputing*, page 9. IEEE Computer Society, 2005.
- [19] Amina Guermouche, Thomas Ropars, Elisabeth Brunet, Marc Snir, and Franck Cappello. Uncoordinated checkpointing without domino effect for send-deterministic MPI applications. In *International Parallel Distributed Processing Symposium (IPDPS)*, pages 989–1000, may 2011.
- [20] Amina Guermouche, Thomas Ropars, Marc Snir, and Franck Cappello. HydEE: Failure containment without event logging for large scale send-deterministic mpi applications. In *IPDPS*, pages 1216–1227. IEEE Computer Society, 2012.
- [21] T. Hoefler, T. Schneider, and A. Lumsdaine. Characterizing the Influence of System Noise on Large-Scale Applications by Simulation. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC'10)*, Nov. 2010.

- [22] T. Hoefler, T. Schneider, and A. Lumsdaine. LogGOPSim - Simulating Large-Scale Applications in the LogGOPS Model. In *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing*, pages 597–604. ACM, Jun. 2010.
- [23] T. Hoefler, C. Siebert, and A. Lumsdaine. Group Operation Assembly Language - A Flexible Way to Express Collective Communication. In *ICPP-2009 - The 38th International Conference on Parallel Processing*. IEEE, Sep. 2009.
- [24] Dewan Ibtesham, Dorian Arnold, Patrick G. Bridges, Kurt B. Ferreira, and Ron Brightwell. On the viability of compression for reducing the overheads of checkpoint/restart-based fault tolerance. *2012 41st International Conference on Parallel Processing*, 0:148–157, 2012.
- [25] Curtis L Janssen, Helgi Adalsteinsson, Scott Cranford, Joseph P Kenny, Ali Pinar, David A Evensky, and Jackson Mayo. A simulator for large-scale parallel computer architectures. *International Journal of Distributed Systems and Technologies (IJDST)*, 1(2):57–73, 2010.
- [26] Sudarsun Kannan, Ada Gavrilovska, Karsten Schwan, and Dejan Milojicic. Optimizing checkpoints using nvm as virtual memory. In *Proceedings of the International Parallel and Distributed Processing Symposium, IPDPS '13, New York, NY, USA, 2013*. ACM.
- [27] Sébastien Monnet, Christine Morin, and Ramamurthy Badrinath. A hierarchical checkpointing protocol for parallel applications in cluster federations. In *Parallel and Distributed Processing Symposium, 2004. Proceedings. 18th International*, page 211. IEEE, 2004.
- [28] Sébastien Monnet, Christine Morin, and Ramamurthy Badrinath. Hybrid checkpointing for parallel applications in cluster federations. In *Cluster Computing and the Grid, 2004. CCGrid 2004. IEEE International Symposium on*, pages 773–782. IEEE, 2004.
- [29] Adam Moody, Greg Bronevetsky, Kathryn Mohror, and Bronis R. de Supinski. Design, modeling, and evaluation of a scalable multi-level checkpointing system. In *ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC '10)*, pages 1–11, 2010.
- [30] Misbah Mubarak, Christopher D Carothers, Robert Ross, and Philip Carns. Modeling a million-node dragonfly network using massively parallel discrete-event simulation. In *High Performance Computing, Networking, Storage and Analysis (SCC), 2012 SC Companion.*, pages 366–376. IEEE, 2012.
- [31] Ron A. Oldfield, Sarala Arunagiri, Patricia J. Teller, Seetharami Seelam, Maria Ruiz Varela, Rolf Riesen, and Philip C. Roth. Modeling the impact of checkpoints on next-generation systems. In *24th IEEE Conference on Mass Storage Systems and Technologies*, pages 30–46, September 2007.

- [32] J. S. Plank, Y. B. Kim, and J. J. Dongarra. Algorithm-based diskless checkpointing for fault tolerant matrix operations. In *Twenty-Fifth International Symposium on Fault-Tolerant Computing. Digest of Papers*, pages 351–360, Pasadena, CA, USA, June 1995. Los Alamitos, CA, USA : IEEE Comput. Soc. Press, 1995.
- [33] James S. Plank, Kai Li, and Michael A. Puening. Diskless checkpointing. *IEEE Transactions on Parallel and Distributed Systems*, 9(10):972–986, October 1998.
- [34] Rolf Riesen, Kurt Ferreira, Jon Stearley, Ron Oldfield, James H Laros III, Kevin Paredetti, Ron Brightwell, et al. Redundant computing for exascale systems. Technical report, Technical report SAND2010-8709, Sandia National Laboratories, 2010.
- [35] Bianca Schroeder and Garth A. Gibson. A large-scale study of failures in high-performance computing systems. In *International Conference on Dependable Systems and Networks (DSN)*, June 2006.
- [36] L. M. Silva and J. G. Silva. An experimental study about diskless checkpointing. In *24th EUROMICRO Conference*, pages 395 – 402, Vasteras, Sweden, August 1998. IEEE Computer Society Press.
- [37] Horst D Simon. Barriers to exascale computing. In *High Performance Computing for Computational Science-VECPAR 2012*, pages 1–3. Springer, 2013.
- [38] T. Hoefler. LogGOPSim - A LogGOPS (LogP, LogGP, LogGPS) Simulator and Simulation Framework. <http://www.unixer.de/research/LogGOPSim/>, Apr. 10 2013.
- [39] Anand Tikotekar, Geoffroy Vallée, Thomas Naughton, Stephen L Scott, and Chokchai Leangsuksun. Evaluation of fault-tolerant policies using simulation. In *Cluster Computing, 2007 IEEE International Conference on*, pages 303–311. IEEE, 2007.
- [40] Patrick Widener, Kurt Ferreira, Scott Levy, Patrick G. Bridges, Dorian Arnold, and Ron Brightwell. Asking the right questions: benchmarking fault-tolerant extreme-scale systems. In *Proc. 6th Workshop on Resiliency in High Performance Computing*, Aachen, Germany, August 2013. in conjunction with Euro-Par 2013.
- [41] Gengbin Zheng, Terry Wilmarth, Praveen Jagadishprasad, and Laxmikant V Kalé. Simulation-based performance prediction for large parallel machines. *International Journal of Parallel Programming*, 33(2-3):183–207, 2005.

DISTRIBUTION:

- 1 MS 1319 Kurt Ferreira , 1423
- 1 MS 0899 Technical Library, 9536 (electronic copy)
- 1 MS 0359 D. Chavez, LDRD Office, 1911

