

Addressing failures in exascale computing

Marc Snir¹, Robert W Wisniewski², Jacob A Abraham³,
Sarita V Adve⁴, Saurabh Bagchi⁵, Pavan Balaji¹, Jim Belak⁶, Pradip Bose⁷,
Franck Cappello¹, Bill Carlson⁸, Andrew A Chien⁹, Paul Coteus⁷,
Nathan A DeBardleben¹⁰, Pedro C Diniz¹¹, Christian Engelmann¹²,
Mattan Erez³, Saverio Fazzari¹³, Al Geist¹², Rinku Gupta¹, Fred Johnson¹⁴,
Sriram Krishnamoorthy¹⁵, Sven Leyffer¹, Dean Liberty¹⁶, Subhasish
Mitra¹⁷, Todd Munson¹, Rob Schreiber¹⁸, Jon Stearley¹⁹ and
Eric Van Hensbergen²⁰

The International Journal of High
Performance Computing Applications
2014, Vol. 28(2) 129–173
© The Author(s) 2014
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1094342014522573
hpc.sagepub.com



Abstract

We present here a report produced by a workshop on ‘Addressing failures in exascale computing’ held in Park City, Utah, 4–11 August 2012. The charter of this workshop was to establish a common taxonomy about resilience across all the levels in a computing system, discuss existing knowledge on resilience across the various hardware and software layers of an exascale system, and build on those results, examining potential solutions from both a hardware and software perspective and focusing on a combined approach.

The workshop brought together participants with expertise in applications, system software, and hardware; they came from industry, government, and academia, and their interests ranged from theory to implementation. The combination allowed broad and comprehensive discussions and led to this document, which summarizes and builds on those discussions.

Keywords

Resilience, fault-tolerance, exascale, extreme-scale computing, high-performance computing

Report on a workshop organized by the Institute for Computing Sciences on 4–11 August 2012 at Park City, Utah.

1 Introduction

‘The problems are solved, not by giving new information, but by arranging what we have known since long.’

– Ludwig Wittgenstein, *Philosophical Investigations*.

This article is the result of the workshop on ‘Addressing failures in exascale computing’ held in Park City, Utah, 4–11 August 2012. The workshop was sponsored by the Institute for Computing in Science (ICiS). More information about ICiS activities can be found at <http://www.icis.anl.gov/about>. The charter of this workshop was to establish a common taxonomy about resilience across all the levels in a computing system; to use that common language in order to discuss existing knowledge on resilience across the various hardware and software layers of an exascale system; and then to build on those results, examining potential solutions from both a hardware and software perspective and focusing on a combined approach.

¹Argonne National Laboratory, IL, USA

²Intel Corporation, CA, USA

³University of Texas at Austin, TX, USA

⁴University of Illinois at Urbana-Champaign, IL, USA

⁵Purdue University, IN, USA

⁶Lawrence Livermore National Laboratory, CA, USA

⁷IBM T.J. Watson Research Center, NY, USA

⁸IDA Center for Computing Sciences, MD, USA

⁹The University of Chicago, IL, USA

¹⁰Los Alamos National Laboratory, NM, USA

¹¹USC Information Sciences Institute, CA, USA

¹²Oak Ridge National Laboratory, TN, USA

¹³Booz Allen Hamilton, VA, USA

¹⁴SAIC, VA, USA

¹⁵Pacific Northwest National Laboratory, WA, USA

¹⁶Advanced Micro Devices, MA, USA

¹⁷Stanford University, CA, USA

¹⁸Hewlett Packard, CA, USA

¹⁹Sandia National Laboratory, NM, USA

²⁰ARM Inc., TX, USA

Corresponding author:

Marc Snir, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue Argonne, IL 60439.
Email: snir@anl.gov

The workshop brought together participants with expertise in applications, system software, and hardware; they came from industry, government, and academia; and their interests ranged from theory to implementation. The combination allowed broad and comprehensive discussions and led to this article, which summarizes and builds on those discussions.

The article is organized as follows. In the rest of the introduction, we define resilience and describe the problem of resilience in the exascale era. In Section 2, we present a consistent framework and terms used in the rest of the document. Sections 3 and 4 describe the sources and rates for hardware and software errors. Section 5 examines classes of software capabilities in preventing, detecting, and recovering from errors. Section 6 takes a system-wide view and describes possible ways of achieving resilience. Section 7 presents possible scenarios and how to handle failures. Section 8 provides suggested actions.

1.1 The problem of resilience at exascale

DOE and other agencies are engaged in an effort to enable exascale supercomputing performance early in the next decade. Extreme-scale computing is essential for progress in many scientific and engineering areas and for national security. However, progress from current top HPC systems (at tens of petaflops peak performance and roughly 1 PF sustained performance) to systems 1000 times more powerful will encounter obstacles. One of the main roadblocks to exascale is the likelihood of much higher error rates, resulting in systems that fail frequently and make little progress in computations or in systems that may return erroneous results. Although such systems might achieve high nominal performance, they would be useless.

Higher error rates will be due to a confluence of many factors:

- Hardware failures are expected to be more frequent (discussed in more detail in Section 3). Errors undetected by hardware may be frequent enough to affect many computations.
- As hardware becomes more complex (heterogeneous cores, deep memory hierarchies, complex topologies, etc.), software will become more complex and hence more error-prone. Failure and energy management also add complexity. In addition, the larger scale will add complexities as more services need to be decentralized, and complex failure modes that are rare and ignored today will become more prevalent.
- Application codes are becoming more complex. Multiphysics and multiscale codes couple an increasingly large number of distinct modules.

Data assimilation, simulation, and analysis are coupled into increasingly complex workflows. Furthermore, the need to reduce communication, tolerate asynchrony, and tolerate failures results in more complex algorithms. The more complex libraries and application codes are more error-prone. Software error rates are discussed in Section 4 in more detail.

1.2 Applicable technologies

The solution to the problem of resilience at exascale will require a synergistic use of multiple hardware and software technologies.

Avoidance: for reducing the occurrence of errors

Detection: for detecting errors as soon as possible after their occurrence

Containment: for limiting the impact of errors

Recovery: for overcoming detected errors

Diagnosis: for identifying the root cause of detected errors

Repair: for repairing or replacing failed components

We discuss potential hardware approaches in Section 3 and potential software solutions to resilience in Section 5.

1.3 The solution domain

The current approach to resilience assumes that silent errors are rare and can be ignored. Applications checkpoint periodically; when an error is detected, system components are either restored to a consistent state or restarted; applications are restarted from the latest checkpoint. We divide the set of possible solutions for resilience at exascale into three categories.

Base option: Use the same approach as today. This would require the least effort in porting current applications but may have a cost in terms of added hardware and added power consumption. We discuss in Section 7.1 what improvements are needed in hardware and system software in order to carry this approach into the exascale range, and we consider what costs will be incurred.

System option: Use a combination of hardware and system software to handle resiliency in a manner that is transparent to the application developer. This approach will require no change in application codes and is therefore equivalent to the base option from the viewpoint of application developers. The relative cost of hardware changes vs. system software changes will dictate preferences between the base option and the system option. We discuss this option in Section 7.2.

Application option: Require application developers to handle resilience as part of their application code. The approach is more invasive from the viewpoint of application developers but may reduce the cost of exascale platforms and their energy consumption. We further subdivide this option into two suboptions.

Application-level error detection: Application code is responsible for error detection; recovery is done, as today, by restarting from a checkpoint. That is, the only added burden on application developers is to provide a checkpoint validation routine. We discuss this option in Section 7.3.

Application-level error correction: Application code is also written so as to avoid the need for global checkpoint and global restart, thus possibly reducing the overheads entailed by this approach. We discuss this option in Section 5.6.4.

We find that some technologies are essential no matter which approach is chosen. For example, it is essential to reduce the frequency of system crashes and to reduce the time to recover from system crashes. Other technologies are ‘no brainers’ in that they improve the resilience of systems with little added cost. This is true, for example, of failure prediction and avoidance, as discussed in Section 5.2.

The three options are not mutually exclusive. The system option will still require adequate hardware support, and the application option will require adequate hardware and system software support. Design choices will need to consider the maturity of various technologies and the relative cost of the different choices of higher platform acquisition cost, higher power consumption, or higher cost for application code development and porting. The balance may change over time and may well not be the same for today’s 10 PF machines as for a 100 PF system or an exascale system. To be able to make the tradeoffs requires understanding the costs based on the expected and possible capabilities at each layer. Thus, we discuss in Section 8 the commonality between these options, pointing out technologies that are clearly needed no matter what path is taken, and the research, observations, and experiments that can help us choose the appropriate path.

1.4 Previous reports

Our work leverages several recent reports on resilience.

A DARPA white paper on system resilience at extreme scale was issued in 2009 (Elnozahy et al., 2009). It points out that high-end systems waste 20% of its computing capacity on failure and recovery. The white paper outlines possible evolutionary and revolutionary research with the goal of bringing this number down to 2%.

Blue Waters and Teragrid co-sponsored a workshop in 2009 on ‘Fault-tolerance for extreme-scale computing’ (Katz et al., 2009). The ensuing report proposes focusing on better communications between vendors, system people, and application teams; more measurements to quantify the problem; and better preventive maintenance.

A DOE/DOD report issued in 2010 (DeBardeleben et al., 2010b) identifies resilience as a major emerging issue for high-end computing (HEC) that requires new approaches. It calls for a national effort and proposes research in five thrust areas: theoretical foundations, enabling infrastructure, fault prediction and detection, monitoring and control, and end-to-end data integrity. This report considers resilience to be ‘concerned with reliability of information in lieu of, or even at the expense of, reliability of the system’.

A report published in 2009 by the NCSA/INRIA Joint Laboratory for Petascale Computing (Cappello et al., 2009) identifies four major research issues for exascale resilience: (1) fault detection, propagation and understanding; (2) fault recovery; (3) fault-oblivious algorithms; and (4) stress testing of the proposed fault-tolerance solutions.

A DOE/DOD report issued in 2012 (Daly et al., 2012) identifies six high priorities: fault characterization, detection, fault-tolerant algorithms, fault-tolerant programming models, fault-tolerant system services, and tools.

The Computing Community Consortium (CCC) organized a ‘Cross-layer reliability visioning study’ in 2011 (DeHon et al., 2011). This study, while not focused on HPC, makes many relevant points. It suggests a research and education program with eight components: repairable hardware architectures; cross-layer information sharing; multilayer error filtering; multilayer tradeoffs for error-handling; differential reliability; techniques, theories, and platforms that are scalable and adaptive to a wide range of error rates and error types; graceful degradation; and embedding of reliability and immunology engineering into electrical engineering, computer engineering, and computer science curricula.

A recent DOE workshop (Geist et al., 2012) focused on resilience from the perspective of DOE, with the following goals: (1) describe the required HPC resilience for critical DOE mission needs; (2) detail what HPC resilience research is already being done at the DOE national laboratories and is expected to be done by industry or other groups; (3) determine what fault management research is a priority for DOE’s Office of Science and NNSA over the next five years; and (4) develop a roadmap for getting the necessary research accomplished.

The International Exascale Software Project Roadmap (Dongarra et al., 2011), which is the result of

more than a year of coordinated effort by a large international group of experts, discusses various aspects of resilience. Resilience is identified as a major cross-cutting issue that requires support at the OS level (APIs for fine-grained management of resilience), run-time (transparent resilience), compilers, applications (application-driven resilience models), and algorithms. The report recommends research on improvements to checkpoint/rollback, and fault-avoiding and fault-oblivious software.

We included in this list only recent reports. We note, however, that research on fault-tolerant computing is as old as computers are. Frequent failures were a major problem in the earliest computers: ENIAC had a mean time to failure (MTTF) of two days (Randall, 2006). Major advances in this area occurred in the 1950s and 1960s, for example in the context of digital telephone switches (Downing et al., 1964) and mainframes (Spainhower and Gregg, 1999). More recently, NASA examined the use of non-rad hardened, commercial-off-the-shelf (COTS) processors for space missions, which requires tolerance of hardware errors (Katz and Some, 2003). Bibliographical research must be an important component of a research program in resilience.

2 Taxonomy of terms

‘Clear language engenders clear thought.’

– Richard Mitchell, *The Underground Grammarian*.

The absence of agreed-upon definitions and metrics for supercomputer reliability, availability, and serviceability has, in the past, obscured meaningful discussion of the issues involved and has hindered their solution (Stearley, 2005). In order to avoid similar confusion, we start by defining our terms. We broadly follow the taxonomy of Avižienis (Avižienis et al., 2004), which has roughly 2000 citations, with additions specific to our domain.

2.1 Dependability

The definitions in this section are based almost entirely on Avižienis et al. (2004).

System: an entity that interacts with other entities

Component/subsystem: a system that is part of a larger system

Atomic component: the point at which system/component recursion stops, by desire or discernability

Functional specification: description of system functionality and performance, defining the threshold between a *correct* and an *incorrect* service (acceptable vs unacceptable)

Service: a system’s externally perceived behavior

Quality of service (QoS): guarantees provided by the system on the performance and reliability of the service it provides

Behavior: what a system does to implement its function, described by a series of states

Total state: a system’s computation, communication, stored information, interconnection, and physical condition

Dependability: the ability to avoid service failures that are more frequent and more severe than is acceptable

Dependence: the extent to which a system’s dependability is affected by another’s

Trust: accepted dependence

The terms ‘fault’, ‘error’, and ‘failure’ are sometimes used synonymously, but we believe that more distinctive use, as defined in Avižienis et al. (2004), is beneficial:

Fault: the cause of an error (e.g. a bug, stuck bit, alpha particle)

Error: the part of total state that *may* lead to a failure (e.g. a bad value)

Failure: a transition to incorrect service (an event, e.g. the start of an unplanned service outage)

Degraded mode/partial failure: the failure of a subset of services

Faults can be *active* or *inactive*, meaning actually causing errors or not. A fault is generally local to a single component, as distinct from errors that may propagate from component to component. Similarly, the failure of one component may lead to the failure of another (i.e. ‘cascading’ failures), as shown in Figure 1.

For example, consider a cracked wire inside a cable. The crack is the fault, and it does not move from cable to cable. Because of the crack, a certain bit may be incorrectly flipped during transmission, resulting in an error (an incorrect bit value). The cable failed to provide correct service. The error may continue to propagate from device to device, perhaps leading to incorrect results (a failure), or that flipped bit may have no effect on final results (no failure).

2.2 Life cycle and operational status

After acceptance, a system is, at any time, in one of the operational states shown in Figure 2.




Figure 1. Error propagation and cascading failures.




Figure 2. System's operational status.

2.3 Failure characteristics

Domain: What has failed. The failure can be involve the wrong *content* (incorrect state) or wrong *timing*: service not provided in a timely manner.

Persistence: A failed system may *halt* (fail-stop) or may exhibit *erratic* behavior.

Detectability: A failure can be *signaled* once it is detected and a warning is generated; otherwise, it is *unsigned*. The detection and signaling mechanism can fail, resulting in *false positives* (a false alarm) or a *false negative* (a failure that did not generate an alarm). The *precision* of a detection mechanism is the fraction of signaled failures that were actual failures, and *recall* is the fraction of failures that were detected and signaled: $precision = 1 - false_positives/signalled$, and $recall = 1 - false_negatives/failures$.

Consistency: A failure is *consistent* if it is perceived identically by all users; it is *inconsistent* (or Byzantine) if it is perceived differently by different users. Fail-stop errors are normally consistent, whereas erratic failures can lead to Byzantine behavior.

2.4 Fault characteristics

Active: fault causes an error

Dormant: fault does not cause an error; the dormant fault is *activated* when it causes an error

Permanent: presence is continuous in time

Transient: presence is temporary

Intermittent: fault is transient and reappears

Hard/solid: activation is systematically reproducible

Soft/elusive: activation is not systematically reproducible

The distinction between hard and soft faults is not a strict one: faults may be due to a complex combination of internal state and external conditions that occur rarely and are difficult to reproduce; they appear as soft faults; a root cause analysis may identify the precise circumstances of the fault, enabling systematic reproduction.

2.5 Error characteristics

Detected: indicated by error message or signal

Latent/silent: not detected

Masked: not causing a failure

Soft: due to a transient fault

2.6 Means of dealing with faults

Forecasting: to estimate the present number, future incidence, and likely consequences of faults

Prevention: to prevent fault occurrence

Removal: to reduce fault number and severity

Tolerance: to avoid service failures in the presence of faults

2.7 Fault-tolerance techniques

Error detection: identify the presence of an error

Concurrent: occurs during service delivery

Preemptive: occurs during planned service outage

Recovery: prevent faults from causing failures

error-handling: eliminate errors

Rollback: revert to previous correct state (e.g. checkpoint, retry)

Rollforward: move forward to a new correct state

Compensation: correct the error (e.g. via redundancy)

Fault handling: prevent faults from reactivating

Diagnosis: identify fault location and type

Isolation: exclude from interaction with other components

Reconfiguration: replace component or move work elsewhere

Reinitialization: perform a pristine reset of state (e.g. reboot)

Error detection identifies the presence of an error but does not necessarily identify which part of the system state is incorrect, and what fault caused this error. By definition, every fault causes an error. Almost always, the fault is detected by detecting the error this fault caused. Therefore, 'fault detection' and 'error detection' are often used synonymously.

'Full diagnosis' identifies the root cause of a failure: the original fault or faults that caused this failure; on the other hand, 'partial diagnosis' traces the error back to previous events in the causality chain but does not necessarily identify the original fault. Thus, failure of a software system may be traced back to a hardware error, such as a bit flip, without identifying the fault that caused this bit flip.

2.8 Metrics

If you can not measure it, you can not improve it.

– Kelvin (1891)

We cannot optimize resilience without measuring it. We discuss two metrics here: workload and availability.

2.9 Workload

A key metric is the ratio of the ideal time to solution on an ideal, fault-free system (T_{solve}) to the actual runtime in a real system ($T_{wallclock}$): *Workload Efficiency* = $T_{solve}/T_{wallclock}$.

In the general case, where the system is running a mix of jobs, we can define workload efficiency as the ratio between the ideal time to solution for this job mix on a fault-free system and the actual running time. The difference between $T_{wallclock}$ and T_{solve} is the *overhead* associated with dealing with faults, errors, and failures, including scheduled downtime, unscheduled downtime, and the cost of detection, diagnosis, repair, compensation, and time lost because of degraded performance.

Typically, workload efficiency is measured with respect to ‘system faults’ and includes all faults underlying applications that impact solution correctness or solution time: software bugs, hardware bugs, hardware faults, and so forth. It does not include faults such as application bugs or user errors. However, the workload efficiency does depend on the application code. For example, it depends on how frequently the user checkpoints and how efficient the checkpoint and restart code are. If failure handling will require increased user involvement in the future, then workload efficiency will increasingly depend on the user code, but the overhead due to user code that handles failures will be increasingly hard to measure.

The workload efficiency metric is an ‘instantaneous metric’. The workload efficiency of a system will vary over time: failure rates are higher on a new system or on a system close to the end of its lifetime. Better system design and better testing procedures may reduce the time needed to stabilize a system and raise the workload efficiency faster. Therefore, it is also useful to define a total workload efficiency metric that integrates workload efficiency over the lifetime of a system. The definition of such an integrated metric has to take into account that computers depreciate rapidly: a flop now is twice as valuable as a flop in two to three years; hence overhead now is twice as expensive as overhead in two

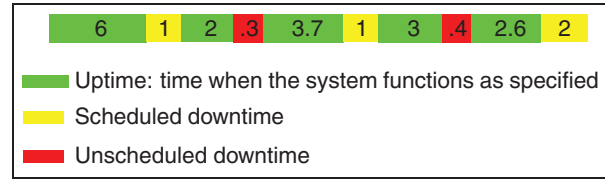


Figure 3. System history.

ratio of the energy needed to solve a problem in an ideal, fault-free system, to the energy needed in reality. Considering the impact of wasted energy is important: some of the techniques for recovery discussed in this report could have little effect on total wall-clock time but could significantly increase power consumption.

In practice, both time and energy are important resources, as are the acquisition cost of the system and the additional program development effort needed to handle failures. The contribution of resilience technology to the value of supercomputers can be measured by a ‘total factor productivity’ (TFP) metric, as the ratio between the cost of inputs (acquisition price, salaries, electricity bills) and the value of outputs (scientific results) (Snir and Bader, 2004). Unfortunately, it is hard to properly estimate the cost of various inputs (e.g. programming time), even harder to separate the contribution of resilience technology from the contribution of other technologies, and practically impossible to put a price on the output of supercomputers.

2.10 Availability

Availability metrics are similar in spirit but more operational in nature. For example, a system may be defined to be ‘down’ when more than 5% of the compute nodes are down or the file system is down; downtime may be considered ‘unscheduled’ if notification occurs less than 12 hours in advance (Mokhtarani et al., 2008).

Consider the time series in Figure 3 of system states, where numbers indicate duration in days.

We tabulate the data into sets and obtain the following statistics:

Set X	$\sum X$	$ X $
Uptime periods={6, 2, 3.7, 3, 2.6}	Uptime = 17.3	NumUptimes = 5
Scheduled downtime periods={1, 1, 2}	Scheduled_Downtime = 4	NumSchedDown = 3
Unscheduled downtime periods={0.3,0.4}	Unscheduled_Downtime = 0.7	NumInterrupts = 2
	Total_Time = 22	

to three years. Given a depreciation rate, it is easy to compute a depreciated total workload efficiency.

The definition of workload efficiency considers time as the critical system resource. If energy is the critical resource, then workload efficiency can be defined as the

The following metric is recommended as a *control* (specified) metric (Bailey et al., 2007): *Scheduled Availability* = $(Total_Time - Scheduled_Downtime) / Total_Time$.

In our example, *Scheduled Availability* = $(22 - 4) / 22 = 81.8\%$.

The following metric is recommended as an *observed* metric: $Actual\ Availability = Uptime/Total_Time$.

In our example, $Actual\ Availability = 17.3/22 = 78.6\%$.

We are using *interrupt* as synonymous with detected failure, so *mean time between interrupts* (MTBI) is equal to *mean time between failures* (MTBF). In our example, $MTBI = TotalTime/NumInterrupts = 22/2 = 11$ days.

Similarly, if MTTI is the *mean time to interrupt*, then $MTTI = Uptime/NumInterrupts = 17.3/2 = 8.65$ days.

The *mean time to repair* (MTTR) is the average length of a unscheduled downtime period. In our example, $MTTR = Unscheduled_Downtime/NumInterrupts = 0.7/2 = 0.35$ days.

The *mean uptime* is the average length of an uptime period. In our example, $MeanUptime = Uptime/NumUptimes = 17.3/5 = 3.46$ days.

The *failures in time* (FIT) rate is the number of failures in a device or system that can be expected in one billion hours of operations. Thus $MTBF = 10^9/FIT$.

2.1.1 Subsystem

When discussing faults, errors, and failures, one must carefully identify what ‘system’ is being referred to. In the previous example, the cable can be considered a system (of wires, solder connections, pins, etc.), the transmission network as a whole can be considered a ‘system’ (of cables, switches, network cards), and the entire collection can be considered a ‘system’ (compute nodes, I/O nodes, network, disks, etc.).

The taxonomy (Avizienis et al., 2004) was developed to address both dependability and security, so the definitions are extremely broad. For example, ‘system’ can refer not only to computing equipment but also to a hacker or group of collaborating hackers. We found it important to identify what is meant by ‘system’ and to identify when that definition changes during the discussion, such as ‘full system’ versus ‘I/O system’. Some uses of ‘system’ include applications, users, and administrators, but the majority of participants referred to ‘full system’ as the collection of components *underlying* the application (not including the application or elements above it, such as users).

Unique acronyms can increase clarity. For example, Sandia and Los Alamos National Laboratories prepend an ‘S’ (e.g. *SMTTI*) to metrics that apply to the full system and other prefixes to identify subsystems (Stearley, 2005). *JMTTI*, the *job mean time to interrupt*, is defined as $JMTTI = (Uptime \times NumJobs)/NumJobInterrupts$, where *NumJobs* is the total number of jobs run and *NumJobInterrupts* is the total number of jobs terminated as a result of any failure. *NMTTI*, *node mean time to interrupt*, is defined as $NMTTI = Uptime \times NumNodes/NumNodeFailures$, where *NumNodes* is the

total number of nodes and *NumNodeFailures* is the total number of node failures.

2.1.2 Statistical models

Analyses of failures and recovery algorithms assume that failures occur according to a probabilistic process that has a closed-form description. A typical assumption is that failures are independent, that is, failure intervals are independent, identically distributed (i.i.d.) random variables. This assumption is clearly false over long periods, since failures are more frequent on a new system or on a system close to the end of its expected lifetime (this leads to a so-called bathtub distribution of failures). It is not clear to what extent the assumption is valid over short time periods, since many phenomena may cause correlated failures. In particular, even if faults are independent, some faults may cause cascading failures of many components. For example, a power or cooling fault can cause the failure of a large number of nodes.

It is often assumed that between-failure intervals have an exponential distribution, with a cumulative distribution function (CDF) $F(t) = 1 - e^{-t/M}$, where M is the MTBI. Such a distribution is implied by the assumption that failures occur according to a Poisson process: the probability that a failure occur during a time interval depends only on the length of this interval. A Weibull distribution, with a CDF of $F(t) = 1 - e^{-(t/M)^k}$, can be used to model a decreasing failure rate ($k < 1$), constant failure rate ($k = 1$), or increasing failure rate ($k > 1$).

An empirical study of HPC failure data from Los Alamos National Laboratory showed a poor fit to an exponential distribution, whereas gamma or Weibull distributions with decreasing failure rates (0.7–0.8) fit well (Schroeder and Gibson, 2010). Surprisingly this study showed that the Weibull distribution fit better in the outer years of the observed system, while no distribution fit well in the first years. These results could be interpreted as meaning that failures in HPC systems are chaotic during the long period it takes for the system to stabilize and that the system keeps improving its reliability through its lifetime. Such an interpretation is consistent with the observation that most failures are due to software. See <http://cfdr.usenix.org> for this and other data.

2.1.3 Resilience, fault tolerance, and dependability

Until now, we have been using the key term ‘resilience’ without clearly defining it. Several reports (Elnozahy et al., 2009; DeBardeleben et al., 2010a,b; Daly et al., 2012) have used different definitions, and debate continues about how, or whether, resilience differs from ‘fault tolerance’ or ‘dependability’. Avizienis et al.

(2004) considered it synonymous with ‘fault tolerance’ and defined it as a wide collection of techniques. The authors defined ‘dependability’ as the ‘ability to avoid services failures that are more frequent and more severe than is acceptable’. In HPC, service failure has two aspects: (1) failure to run a program, or incorrect answer and (2) computation taking too long. The second criterion is quantitative and can be measured in various ways, in particular by using the workload efficiency metric defined earlier: a system fails if its workload efficiency is below a certain threshold. Accordingly, *resilience* can be defined as follows:

The collection of techniques for keeping applications running to a correct solution in a timely and efficient manner despite underlying system faults.

‘Correct’, ‘timely’, and ‘efficient’ are context-dependent. In some contexts ‘correct’ may mean ‘bit-reproducible’; in another context, it could mean ‘within a rounding error’; in yet another context, we could be content with a system that frequently provides a correct solution to a problem, provided that we can efficiently verify solutions. ‘Timely’ and ‘efficient’ are relative rather than absolute (as in before the hurricane arrives and within our power budget). The definition of ‘efficient’ also depends on what we consider to be the total system; for example, are programming costs included?

3 Sources and rates of hardware faults and errors

In this section we describe a generic HPC machine along with the various hardware errors and failures that can occur while it is executing an application. We focus on hardware aspects and do not account for any masking or handling in software. We summarize the

rates at which these errors and component failures occur on current systems and then discuss models for the underlying fault mechanisms, project these models to future 11 nm technology, and recommend possible mitigation techniques and their overheads.

3.1 Generic machine model and associated errors and failures

Figure 4 describes a generic exascale machine, patterned after the current generation of HPC machines at Argonne, Los Alamos, Lawrence Livermore, and Oak Ridge National Laboratories and similar leading supercomputing centers. Faults can occur in any part of the machine, with differing consequences. Some failures (fans, power converters) are masked by redundant hardware. Other failures (nodes) will cause an application to crash and restart from the last checkpoint with a new set of nodes but will not cause the system to crash. Some failures cause the entire system to crash and have to be rebooted. The severity of different failures can be measured by the loss of machine time they cause. The masked failure of a fan slightly increases scheduled downtime; a system crash causes the entire machine to be down for half an hour or more.

To expand on the hierarchy, we imagine that the nodes, servers, and switches of the machine are composed of field replaceable units (FRUs): processors, memory modules, various circuit cards, power and fan modules, and the like, which are usually collected into removable and field serviceable drawers. Sets of drawers may form chassis, and multiple chassis form racks. Typically, but not always, communication is highest between FRUs on a processor node (formed of one or more processor sockets sharing coherent memory and

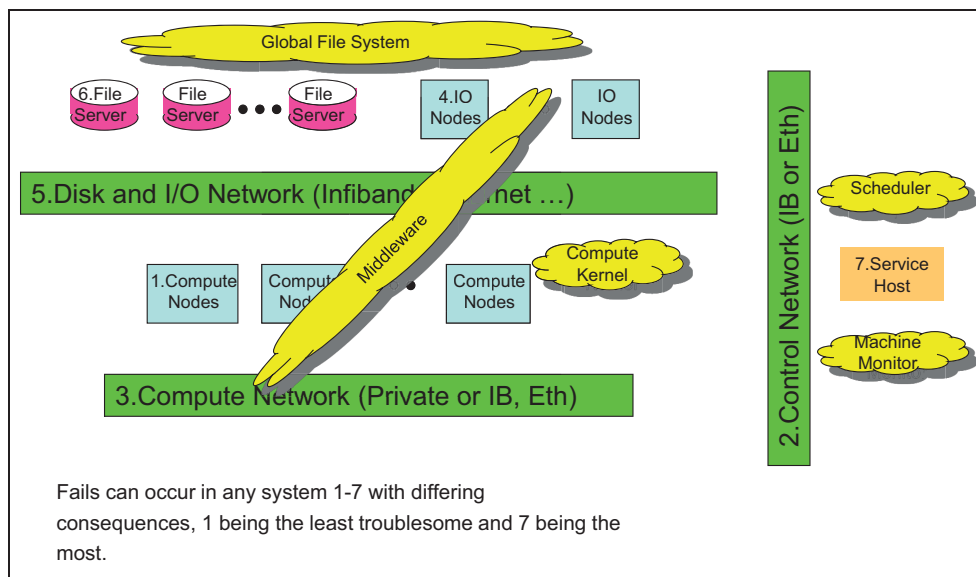


Figure 4. Generic exascale machine.

with at least one network connection). This then is a natural fault containment region. Further, groups of nodes may share some common resource (a network adapter card, power supply, or fan module) making this group of nodes in a drawer (or chassis) a higher-level containment region. The entire rack, perhaps sharing a common resource such as a power cord and bulk AC-DC power supply, could form an even higher containment region. However, HPC applications are tightly coupled, so that errors propagate quickly across components. Software help is needed in order to avoid error propagation and transform physical fault containment units into logical error containment units.

3.2 Classification of errors and failures

Hardware faults can result in errors and failures that may be grouped into three categories: (1) detected and corrected by hardware (DCE), (2) detected in hardware but flagged as being uncorrectable (DUE), and (3) silent (SE). A silent fault may be masked; an SDC is an error caused by an unmasked silent fault. We describe these categories below and discuss the possibility that faults may lead to operating with degraded performance, efficiency, and/or fault protection capability.

Examples of DCE: (a) a detected error in an error checking and correcting (ECC)-protected random access memory (RAM) array that is corrected ‘in place’ before being passed on to a unit that consumes that piece of data and (b) a detected parity error in the processor pipeline that triggers an instruction retry mechanism, resulting in recovery of an uncorrupted, prior-architected register state and re-execution from that point. In the latter case, the recovery mechanism must ensure that leakage of potentially corrupted data to the system’s memory or I/O state is prevented during the whole ‘detection and recovery’ process. The system can be architected such that DCEs are usually transparent to the user (application) program and possibly even to supervisory system software (e.g. OS). In some cases the supervisory system or OS is invoked in order to help record DCE statistics in system memory for later error analysis. In such cases, the DCE is still transparent to the user application. Usually, hardware has autonomous (software-transparent) mechanisms to record DCE statistics in hardware trace (debug) arrays for later diagnostics. Note that frequent DCEs will slow the system and could, in extreme cases, cause timing errors.

Examples of DUE: (a) a double-bit error, detected during the attempted reading of a SECDED ECC-protected RAM array datum, that could not be corrected ‘in place’ and (b) a detected parity error in the processor pipeline that cannot inform the on-chip recovery mechanism within a stipulated deadline, which is an architected parameter designed to ensure that

known (potentially) corrupted data is not released to system memory or I/O state. Usually, all DUEs are flagged as an exception to system software by the hardware. Depending on the nature (severity) of the DUE and the capability of the system, the software should be able to handle the hardware-raised exception in a manner that enables one of the following three actions: (i) restart of the processor execution from a local or global checkpoint; (ii) application checkstop that terminates the application, without crashing the node; or (iii) system checkstop that results in a machine check (requiring ‘reboot’) for the particular node or, in the worst case, perhaps even the whole system. In some cases it might be preferable to simply mark corrupted values as invalid, or poisoned, and allow the application itself to handle the error. An example is to use NaN values to prevent incorrect data from silently corrupting results, while still allowing for potential application-level masking or handling.

Examples of SDC: (a) an undetected arithmetic computation error, within an integer or floating-point data path pipeline, that makes it into architected register state (and eventually perhaps the system memory state) without triggering any error alert at the hardware level; (b) an undetected control error that results in a premature termination of an iterative loop computation that may result in a datum held in register or memory state to contain a value that is incorrect from a program-intended perspective; and (c) incorrect memory and network transfers that were not detected by the error protection mechanisms (e.g. triple-bit errors with SECDED protection). Such SDCs may eventually be detected within a self-checking application program or as a result of a triggered DUE, but such a detection could happen many thousands, millions, or billions of cycles beyond the point of the original occurrence of the SDC. Thus, a sophisticated ‘root cause analysis’ of a DUE may later point to an originating (causative) SDC when it comes to proper accounting statistics of various categories of errors in the hardware.

As a consequence of errors originating from hardware sources, and the associated error-handling hierarchy in hardware and/or software, the overall computing system may manifest degraded levels of QoS as viewed by the end user. For example, if the system encounters a node failure, even if the system or application can recover from the failure, the system will operate at a degraded performance level during the period of system reconfiguration (via updates in the routing tables, etc.). Similarly, an escalated sequence of ECC memory errors may eventually result in a memory ‘chipkill’ that reduces the amount of available system memory (before the defective memory module is replaced), thereby degrading performance. Similarly, certain other repair actions resulting from the flagging of hardware errors

Table 1. Error and failure rates for the Intrepid Blue Gene/P system.

Detected uncorrectable	Predicted % fails per repair period	Intrepid (ANL 40 racks) observed failures per repair period	Intrepid without I/O failures per repair period
Compute cards	90%	0.648	0.648 (72%)
Node boards	5%	0.137	0.137 (15%)
I/O cards	2%	0.785	0.000 (0%)
Link cards	2%	0.020	0.020 (2%)
Service cards	1%	0.098	0.098 (11%)
Fans	0%	0.000	0.000 (0%)
Bulk power	0%	0.000	0.000 (0%)
Mid-planes	0%	0.000	0.000 (0%)
Clock card	0%	0.000	0.000 (0%)
		1.69	0.90
Detected and corrected/marked			
Compute cards (80% DRAM)	58%	2.003	2.003 (53%)
Node boards	28%	0.491	0.491 (13%)
IO cards	0%	0.000	0.000 (0%)
Link cards	2%	0.059	0.059 (2%)
Service cards	1%	0.196	0.196 (5%)
Fans	4%	0.079	0.079 (2%)
Bulk power	6%	0.884	0.884 (24%)
Mid-planes	0%	0.000	0.000 (0%)
Clock card	0%	0.000	0.000 (0%)
		3.71	3.71

may reduce the capability of hardware in terms of being able to detect the full range of errors that the system was originally designed for.

3.3 Quantification of component errors and failures

Table 1 shows the hardware error and failure data for 382 days of the Intrepid system at Argonne National Laboratory. This 40-rack, 557 TF Blue Gene/P system currently shows a mean time to (hardware) interrupt of 7.5 days. Thus, the total of any detected hardware failure, including compute nodes, I/O nodes, compute node interconnect, control hosts, and file servers, was roughly one per 7.5 days. This was extremely close to the seven-day MTBF predicted for the machine back in 2006, well before installation. We point out that this agreement was obtained only after wholesale replacement of two minor but problematic elements of the machine: the 10 Gb/s optical transceiver modules on the I/O links and early versions of the bulk power supply modules. This experience is consistent with the Los Alamos National Laboratory study discussed in Section 2.12: in the beginning there is chaos; statistical regularity takes over when the system matures. Also, while the system failure rate was predicted within 10%, the relative contribution of the different components was quite different from predicted; errors canceled each other.

Not all failures have the same impact. A node board failure affects all 32 compute cards sitting on it (each

card contains a four-core processor and attached memory). The failure of an I/O card can affect all compute cards on the board containing the I/O card. The failure of a link card affects an entire partition or set of nodes that are assigned to a running job.

3.4 Hardware fault, error, and failure models and projections

To project the hardware error and failure rates expected in an exascale machine, one must understand the root cause of these events. While reasonably good models exist for some faults in some components, important gaps remain in the projections we will be able to make. We summarize our best-effort models below.

3.4.1 Compute node soft errors. Soft errors in the compute node (processor and memory only; network, power, and cooling are discussed later in this section) are most often a result of events that are entirely external to the system and cannot be replicated. By far the most significant source of transient faults is energetic particles that interact with the silicon substrate and either flip the state of a storage element or disrupt the operation of a combinational logic circuit. The two common sources of particle strike faults are alpha particles that originate within the package and high-energy neutrons. Alpha particles are charged and may directly create electron-hole pairs. When a high-energy neutron interacts with the silicon die, it creates a stream of secondary charged

Table 2. Summary of assumptions on the components of a 45 nm node and estimates of scaling to 11 nm.

	45 nm	11 nm
Cores	8	128
Scattered latches per core	200,000	200,000
Scattered latches out of cores	632,000	2,530,000
FIT per latch	10^{-1}	10^{-1}
Arrays per core (MB)	1	1
FIT per SRAM cell	10^{-4}	10^{-4}
Logic FIT/latch FIT	0.1–0.5	0.1–0.5
DRAM FIT (per node)	50	50

particles. These charged particles then further interact with the semi-conductor material, freeing electron-hole pairs. If the charged particle creates the electron-hole pairs within the active region of a transistor, a current pulse is formed. This current pulse can directly change the state of a storage device or can manifest as a wrong value at the end of a combinational logic chain.

To analyze the impact a particle strike has on a compute node, we model the effect on each node component separately, namely static RAM (SRAM), latches, combinational logic, dynamic RAM (DRAM), and non volatile RAM (NVRAM). We then determine a rough estimate for the number of units of each component within the node. We use this estimate to provide a rough order-of-magnitude fault rate for the compute node. We also briefly mention how such faults are handled in processors today, and we discuss how advances in process technology are expected to affect these soft faults. We make projections for the impact of particle-strike soft errors on a future 11 nm node, as well as present an estimate of the overhead/error-rate tradeoffs at the hardware level. The estimates are based on the models below and on some assumptions about the components of a node, as shown in Table 2. First, however, we give a few important caveats about the models and projections:

- The numbers summarized in Table 2 do not include errors due to hard faults or to transient faults other than particle strikes. We expect those to be a significant contributor to software-visible errors and failures.
- We do not have access to good models for the susceptibility of near-threshold circuits and do not consider such designs.
- We give only a rough, order-of-magnitude (at best) estimate; many important factors remain unknown with respect to a 11 nm technology node.

We expect that, over the next few years, ongoing research at microelectronic companies, research labs, and in academia will provide more accurate estimates.

We estimate the number of scattered latches out of core as $scattered_latches_per_core \times \sqrt{n_{cores}} \times 1.25$.

SRAM. Large SRAM arrays dominate the raw particle-strike fault rate of a processor silicon die. When a particle strike releases charge within an active region of a transistor in an SRAM cell, the charge collected may exceed the charge required to change the value stored in the cell, causing a single event upset (SEU). An SEU may impact a single SRAM cell or may change the values of multiple adjacent cells. Such multicell upsets (MCUs) are also called *burst errors*. A reasonable ballpark figure for SRAM particle-strike upset rate is one upset every 10^7 hours for 1 Mb of capacity, which is a rate of 10^{-4} FIT/bit (Slayman, 2011). Our best estimates indicate that the SEU rate for SRAM will remain roughly constant as technology scales. While many complex phenomena impact susceptibility, the current roadmap of changes to devices, operating voltage, and scale do not point to extreme changes in susceptibility. What is expected to change is the distribution of MCUs, with a single upset more likely to affect a larger number of cells at smaller scales.

Because the raw FIT/chip from SRAM is high (estimated at roughly 0.5 upsets per year per chip, or multiple upsets an hour in a large-scale HPC system), large arrays are protected with error detection and error correction capabilities. An approach in use today is a combination of physical word interleaving coupled with an error-detection code or with ECC mechanisms. Given the distribution of MCUs today, four-way interleaving with SECDED capabilities per array line is sufficient. Stronger capabilities will likely be needed in the future, but their energy and area overheads are expected to be low (see Table 3). Note that our estimates assume that 4-bit or larger bursts increase from 1% of all SEUs to 10% or higher between 45 nm and 11 nm technology and that the rate of bursts of 8 bits or larger increases from 0.01% of all SEUs to 1% of all SEUs (Ibe et al., 2010).

Note that alternative storage technology with much lower particle-strike error rates is possible. Some current processors use embedded DRAM for large arrays, and future processors may use on-chip arrays of non-volatile storage. Embedded DRAM has an error rate 100 times or more lower than does SRAM. Nonvolatile storage cells are immune to particle strikes but do display some soft-error fault mechanisms (see discussion below).

Latches. The error mechanisms and trends for latches are similar to those of SRAM, and the per-latch SEU rate is expected to remain roughly 10^{-4} – 10^{-3} FIT/bit (Dixit et al., 2009). Given the smaller number of latch cells in a processor today compared with SRAM cells, the overall contribution to error rate of latches is much smaller as well. Future processors will contain a much larger number of latch cells, and protection may be necessary. The protection mechanisms and overheads of latches depend

Table 3. Summary of per-processor particle-strike soft-error characteristics within a compute node (sea level, equator). Note that other sources of transient faults cannot be ignored.

	Array interleaving and SECEDED (Baseline)					
	DCE [FIT]		DUE [FIT]		SE [FIT]	
	45 nm	11 nm	45 nm	11 nm	45 nm	11 nm
Arrays	5000	100,000	50	20,000	1	1000
Scattered latches	200	4000	N/A	N/A	20	400
Combinational logic	20	400	N/A	N/A	0	4
DRAM	50	50	0.5	0.5	0.005	0.005
Total	1000–5000	100,000	10–100	5000–20,000	10–50	500–5000
	Array interleaving and SECEDED (11 nm overhead: ~1% area and <5% power)					
	DCE [FIT]		DUE [FIT]		SE [FIT]	
	45 nm	11 nm	45 nm	11 nm	45 nm	11 nm
Arrays	5000	100,000	50	1000	1	5
Scattered latches	200	4000	N/A	N/A	20	400
Combinational logic	20	400	N/A	N/A	0.2	5
DRAM	50	50	0.5	0.5	0.005	0.005
Total	1500–6500	100,000	10–50	500–5000	10–50	100–500
	Array interleaving and SECEDED + latch parity (45 nm overhead ~10%; 11 nm overhead: ~20% area and ~25% power)					
	DCE [FIT]		DUE [FIT]		SE [FIT]	
	45 nm	11 nm	45 nm	11 nm	45 nm	11 nm
Arrays	5000	100,000	50	1000	1	5
Scattered latches	200	4000	20	400	0.01	0.5
Combinational logic	20	400	N/A	N/A	0.2	5
DRAM	0	0	0.1	0.0	0.1	0.001
Total	1500–6500	100,000	25–100	2000–10,000	1	5–20

on how the latch is used. Some latches are organized in arrays, like SRAM arrays, while other latches are scattered within logic blocks. Array latches can be protected with interleaving and ECC, although such latches are often accessed with finer granularity than large SRAM arrays, which increases the relative overhead of protection. We include this extra cost in Table 3 and project a higher power overhead than area overhead for protecting arrays in order to account for the added protection of latch arrays that may be needed in future processors.

‘Scattered latches’ are more difficult to protect, on the one hand, because the overhead of interleaving and ECC is exorbitant without the regularity of an array. On the other hand, an error in a scattered latch is often masked by the natural operation of the circuit it is part of. Various estimates exist for the derating factor that should be applied for this natural masking, typically ranging from 90% to 95%. The masking rate may depend on the application and also on the architecture, with more streamlined architecture potentially having a lower rate of masked latch errors. If needed, scattered

latches can be protected against particle-strike-induced upsets. The two main techniques that can be applied are hardened latches or a combination of parity prediction from logic with parity checking on a collection of latch bits. Both techniques can be effective but potentially have high overheads if a large fraction of latches must be protected. We show the impact of this overhead in Table 3.

Combinational logic. The trends we expect for particle-strike-induced soft errors in combinational logic are again consistent with those for SRAM and latches. The raw SEU rate associated with combinational logic can reasonably be estimated at 0.1–0.5 FIT for every 1 FIT contributed by scattered latches within logic blocks (Gill et al., 2009). Note that this is the raw upset rate and does not account for logical masking effects. Similar to latches, even if an output of a logic gate is changed, this change is highly unlikely to impact the final result of the circuit. Because the output of a combinational logic path is always a latch, the overall

masking rate of combinational logic upsets is most likely close to 99%.

Note that the raw upset rate quoted above already accounts for electrical masking, which results from the SEU current pulse being attenuated as it passes through multiple gates, and for timing or latch masking, which results from the output of the combinational logic being observed for only a fraction of the cycle. As with scattered latches, we expect the raw fault rate to stay roughly constant as technology scales, and application and architecture may impact masking rates. The parity-prediction mechanism that can be used to detect errors in latches will also detect a large fraction of logic errors. Other techniques for detecting combinational logic soft faults at the hardware level include those based on arithmetic coding (Avizienis, 1973; Rao, 1974; Lo et al., 1989; Lo, 1994) and replication (Austin, 1999; Slegel et al., 1999; Saxena and McCluskey, 2002). Moreover, electrical masking can be increased by using less area and power-efficient gate designs (Hazucha et al., 2003; Lunardini et al., 2004).

DRAM. As a rule, DRAM exhibits a fixed rate of particle-strike soft errors per DRAM die, regardless of technology. This rate is roughly 10–20 FIT/device, and a significant fraction affects multiple bits and entire rows, columns, or banks of the DRAM device (Sridharan and Liberty, 2012). Many DRAM devices are required for the capacity of each node. Recent studies have shown that the error rate of DRAM is far higher than the particle-strike soft errors, indicating that hard faults in either the peripheral or signaling circuits are the main cause of problems (Schroeder et al., 2009; Hwang et al., 2012; Sridharan and Liberty, 2012).

Regardless of the fault mechanism, DRAM is protected with ECC, with large-scale systems typically supporting some form of chipkill-level ECC, which is effective against hard errors as well. We expect that even if new ECC schemes are needed in the future, their overhead will overall be similar to the overhead observed today for most applications.

NVRAM. We cover several technologies: NAND Flash, spin-transfer torque (STT) memory, phase-change memory (PCM), and resistive memories such as *memristor*.

NAND Flash is vulnerable to soft errors. The FIT rate per bit is growing with process shrinks. Currently it is 10^{-5} FIT. It was 10^{-8} FIT at the 100 nm technology node. ECC is needed and used to cope with this rate, which exceeds that of DRAM. NAND Flash wears out after approximately 10^6 rewrite cycles. Many architectural techniques are used to spread the load across the cells of a chip, a technique called wear leveling. Wear-out is not a major issue in consumer storage devices such as media cards. It may be an issue in solid-state disks, but it is clearly manageable there. As main

memory and cache, Flash is unsuitable for this and other reasons.

SST, the leading magnetoresistive random-access memory (MRAM) technology, is under development by Toshiba and Hynix, which have made prototypes at 30 nm. Samsung has made a device at 17 nm. It is dense ($6F^2$ feature size). Speed and energy cost are good. Chips of 1 Gb are under development and may reach the market in 2014. Wear-out does not appear to be a concern for STT. It also seems that STT bits cannot be flipped by particle strikes. Thermal noise seems to cause something similar to soft errors: errors due to external stimuli, not internal imperfections. A FIT rate of 10^{-10} FIT/bit has been reported. Hard errors are an issue. It is said that ‘imperfections in the fabrication process greatly affect the reliability of data in STT-MRAM. Process variability causes variation in the tunneling oxide thickness and cross-section area, which affects both the static and dynamic behaviors of magnetic tunnel junctions, resulting in cell errors’ (Cal 2013). Appropriate responses could include testing and map-around for bad cells, spare cells, and ECC. (These comments likely apply to all the memory technologies we consider.)

PCM is resistive memory in which the state of a chalcogenide glass is changed between crystalline and amorphous by heating and either slow or fast cooling. The resulting change in the electrical resistance determines the state. Multilevel cells are possible with perhaps two bits per cell, but possibly fewer (as when three resistance levels are used). Micron is marketing 45 nm PCM for consumer applications today. PCM has better endurance than Flash, but it may wear out after as few as 10^6 up to a high estimate of 10^9 cycles because of the physical stresses of repeated heating and cooling. It appears to be invulnerable to particle-induced soft errors. The resistance of the PCM cell changes with time. Thermal disturbance due to the heating required for reset of a nearby cell is a chief cause of resistance drift, and this limits cell density. The decay of the stored data is similar to the charge leakage in the DRAM capacitor and, like it, may cause errors. Some combinations of refresh and ECC can cope with drift. Because of the necessity for refresh to arrest drift, it is not clear that PCM is as nonvolatile as necessary for use in offline storage. The rate of required refresh will depend on the degree to which storage density is boosted by using multilevel cells; the tighter the level spacing, the more frequently the cell must be refreshed. In particular, a 1-bit cell with only two levels would have no drift problem. There is thus a complex design space in which density, the cost of mitigating the resistance drift, the data retention time, and the error rate are in competition. Optimization of the PCM cell and its required refresh and error correction architecture is an area of ongoing research.

The memristor is a new technology under development at HP and Hynix and in other laboratories and companies. A memristor stores information in the resistance of a cell (like PCM), that resistance being a function of the past flux (the integral of current) that has passed through (not by heating, as in PCM). HP and Hynix have explored memristive systems in which a metal oxide (often titanium oxide but recently also oxides of tantalum, zirconium, and hafnium) sandwiched between electrodes is electrochemically changed by the passage of current. Memristors are two terminal devices. Like other resistive random-access memory, memristors appear invulnerable to particle-strike soft errors, but a nonrecurring or transient error mechanism in memristors may exist. Recent experimental studies and models show a tendency to fail to remain in the lowest resistance state, randomly, with a probability that is strongly temperature-related. At 150°C, half of all cells may fail if left unchecked for 10 days (see for example Gao et al., 2011; Yu et al., 2012). It is not clear whether these errors are due to cell deficiencies and can be reduced by mapping out bad cells or are totally random and need to be handled by ECC and scrubbing, or both. Memristors can wear out, but the wear-out mechanisms are not as clearly understood as they are for PCM. The ultimate durability of memristors is still to be determined. In new work on tantalum oxide memristors, endurance of over 10 billion cycles has been demonstrated (Yang et al., 2010). Hard-error vulnerabilities appear to be due to wear-out, manufacturing issues, and interface/communication issues and may be comparable to those of PCM.

Nonvolatile memory as a resilience enhancer.

Nonvolatile memory (NVM) is often less vulnerable to soft error due to cosmic rays than is DRAM, but this is almost totally irrelevant to our discussion (since other errors predominate, and NVM has its own sources of errors). Thus, replacing DRAM with NVM will not, in and of itself, enhance resilience. Checkpointing, normally at the application level, is the current default for preserving the state of an ongoing computation in order to protect it from a subsequent failure. Because of the growing size of application state and the failure of disk-based file systems to provide proportionally growing bandwidth, checkpointing to shared disk is not seen as a sustainable approach at exascale. We expect that on-node NVM will appear, for many reasons. One reason is to serve as fast checkpoint storage, since write bandwidth will be superior to disk. In order to cope with node hardware failure, the checkpoint NVM may be ‘twin-tailed’ (capable of being read by a service node or another compute node following node failure). Alternatively, the checkpoints may need to be delocalized, stored on a buddy node, or made recoverable by another scheme. DRAM can also be used for

delocalized checkpoints. It will survive node failures but not global power failures. Since such DRAM will be on standby mode most of the time, there is no significant difference in power consumption.

NVM may serve other resilience functions, in part simply by providing enough memory to do more or as the top level in a hierarchy of nonvolatile storage components. For example, it can be used for logging messages, in order to support local, uncoordinated checkpointing, or for holding file system caches.

3.4.2 Compute node hard errors and failures. While we could provide rough quantitative projections of particle-strike-induced soft-error rates, we cannot ignore possible failures and errors (detected and undetected) due to hard faults. Because of the complexity of designing and efficiently operating future processors, some failures and errors may be intermittent and manifest only with certain environmental conditions or specific execution characteristics. Major concerns include increased early-life failure rate, permanent and intermittent faults associated with device degradation, and increased storage element error rates because of low-voltage operation. Quantitative data on how such hard-fault sources will evolve over technology generations is difficult to predict. But the effects can be enormous. We briefly discuss the issues below.

Early-life failures (infant mortality). Burn-in for screening early-life failures is becoming increasingly challenging (Nigh and Gattiker, 2000; Kundu et al., 2004; Borkar, 2005; Carulli and Anderson, 2005; Van Horn, 2005). Major challenges include power dissipation, cost, and possibly reduced effectiveness and coverage of the burn-in test techniques. Burn-in alternatives, for example Iddq testing (measuring the supply current, or Idd, in the quiescent state) and very low voltage testing (Hao and McCluskey, 1993; Gattiker et al., 1996; Chan et al., 1998; Maxwell et al., 2000), are also experiencing limitations, including high leakage, process variations, and reduced voltage margins. At a highly scaled technology node with minimal reliance on burn-in, the effects of early-life failures can be significant: on the order of several thousands of defective parts per million. Such a high rate of failures is roughly equivalent to adding 10^3 – 10^4 FIT to the node failure rate. More aggressive online techniques for detecting these failures may become necessary.

Device degradation (aging). Device degradation induced by degradation mechanisms such as bias temperature instability (BTI) (Agostinelli et al., 2005; Reddy et al., 2005; Zhou et al., 2010), hot-carrier injection, time-dependent dielectric breakdown, or metal electromigration is becoming important. While design margins (guard bands) are being squeezed to achieve higher

energy efficiency, expanded design margins are required to cope with aging. Hence, traditional speed or voltage margins to overcome degradation may become too expensive. Some projections predict that beyond the 14 nm technology node, guard bands due to BTI degradation may grow to 20% or more, degrading efficiency and performance by a similar amount. Such guard bands are highly dependent on the workload, and quantitative projections can be highly pessimistic for worst-case workloads. Moreover, for near-threshold voltages of operation, a huge dilemma arises: while low-voltage operation can reduce the amount of aging, high-voltage turbo modes of operation or fast execution followed by low-voltage operation for energy efficiency can significantly exacerbate this aging effect. Techniques that dynamically adjust guard bands and improve performance and efficiency have been suggested, but their impact on intermittent failures and errors has not been fully evaluated. Here, the difference between exascale and commodity small-scale systems is vast, because of the scale multiplier of base rates and the impact of large variances on tightly coupled systems.

Low-voltage storage-element stability. As supply voltage is reduced in order to improve energy efficiency and reduce power consumption, maintaining the integrity of storage elements, including latches, flip-flops, and SRAM cells, is challenging. For example, $V_{cc_{min}}$ -related errors can induce so-called Goldilocks failures (Nassif et al., 2012): failures that appear hard but are, in fact, caused by phenomena typically associated with soft failures. Such failures are expected to become more problematic with increasingly complex circuits and lower voltage supplies, affecting circuit structures besides SRAM. At present, the only viable way to deal with $V_{cc_{min}}$ errors in sequential elements is to rely on (expensive) circuit-design techniques or resort to high-voltage operation, resulting in poor energy efficiency.

Possible mitigation techniques. Understanding the effects of such failures is not enough. The question is, how do we mitigate them, especially for silent errors that may lead to SDC? Techniques in the literature that can be useful include (1) online self-test and diagnostics, (2) concurrent error-detection techniques (similar to soft errors), (3) adaptive self-tuning and online optimization, and (4) online self-repair. However, these techniques are generally not supported extensively for existing processors. If the U.S. Department of Energy has to rely on COTS components, chances of all these techniques being supported get even lower. That brings up the question ‘what hardware and software support is required for future exascale systems?’

3.4.3 Network. The transport layer of the network, whether electrical or optical, can be instrumented for

error detection and correction with quantifiable cost. Thus, for example, on Blue Gene/Q, a combination of CRC, Reed–Solomon codes, and Hamming codes, along with a retry mechanism for detected but uncorrected errors, reduces the possibility of an error escape to 10^{50} (Chen et al., 2011). Thus, network transport errors are containable. Network logic, on the other hand, comprises SRAM, latches, and logic, as described above, with their failure modes and correction techniques. Errors that result in data sent to a wrong destination are potentially the most damaging but may be mitigated with hardware or software techniques that use knowledge of the desired and actual recipients to trap errors before data corruption occurs. Networks that support superior error detection and correction, with tailored mechanisms to ensure correct delivery, will surely be a part of exascale systems.

3.4.4 I/O. The increased density of disks results in increased error rates, including an increase in undetected disk errors: those that are not detected by current techniques (RAID 6 included). Various techniques are available for detecting such errors and correcting them, mostly in the form of added redundancy (Hafner et al., 2008). In addition, disk failure rates are often higher than the nominal MTBF would indicate, with a 2%–4% yearly failure rate common (Schroeder and Gibson, 2007); one parity block (RAID 5) is not sufficient, since the probability of two disk failures within the same group is too high. The IBM GPFS system implements in software a RAID 6 scheme (two parity blocks) that can overcome two disk failures (Fadden, 2012).

While these techniques can practically eliminate the risk of data loss, they come at a cost: the disk storage system of a large supercomputer will have continuous I/O background activity due to RAID reconstruction after disk failure. The problem is worsened by the increasing gap between disk capacity and disk bandwidth, which results in increasing reconstruction time, or the need to spread reconstruction across more disks. This background activity will reduce the effective I/O bandwidth and cause significant I/O performance jitter.

3.5 Commercial trends

The technology analysis in this section provides insight into the cost of producing components with acceptably low failure rates; it does not tell us what the price of processors that incorporate these technologies will be. While predicting component prices a decade ahead may be infeasible, we point out that market trends are not favorable. High levels of resilience are important for high-end servers, such as mainframes or RISC/Unix servers (Sun, Power, Itanium). For many other markets (mobile, clouds) vendors are likely to accept lower

reliability in order to achieve lower cost and lower energy consumption. Unfortunately, the market for high-end servers is currently shrinking; the decline is particularly sharp for high-end RISC Unix servers. While some of this decline may be attributed to the current state of the economy, this sector clearly is an increasingly small fraction of the IT industry. Furthermore, this sector is likely to be less price-sensitive than other sectors. Buyers of mainframes or high-end Unix servers have been willing to accept large markups on price per performance, in order to achieve higher reliability levels. They are also less sensitive to power consumption, both because they are less sensitive to operation cost, and because high-end servers usually are a small fraction of the IT infrastructure. These trends are likely to lead to an increasing cost differential between low-reliability components and high-reliability components and to an absence of high-reliability, low-power components. Systems built with high-end RISC processors (Sparc64, Power7) are already rare in the Top500 list.

3.6 Shielding

The impact of particle strike can be reduced by shielding (an area where DOE has significant expertise). The atmosphere is a natural shield, with higher locations suffering from higher strike rates; a computer at sea level will fail less frequently than one at a high-altitude location. Natural or artificial shielding can further reduce the neutron flux. For example, 2 m of concrete will reduce the impact of 10 MeV neutron radiation by three orders of magnitude (Seltborg et al., 2005); less energetic neutrons are attenuated much more. On the other hand, neutrons with energies above 10 MeV carry a very small fraction of the total energy of cosmic-ray neutrons (Hess et al., 1959). Hence, the cheapest way of avoiding the effect of cosmic radiation-induced errors may be to locate future exascale systems in abandoned tunnels of the defunct superconducting super collider or repurposed atomic shelters.

4 Sources and rates of software faults and errors

A large fraction of system failures is due to software, rather than hardware. A study of major DOE supercomputers in 2004–2006 showed that about 65% of failures could be attributed to software (Oliner and Stearley, 2007), whereas a study of failures in 2012 on Intrepid, the BG/P system at Argonne National Laboratory, showed that less than 16% of job crashes were due to hardware problems (Allcock, 2013, private communication). Results of studies show variance, however; a study by Schroeder and Gibson in 2010 showed that failures attributable to hardware ranged from 30%–60% (Schroeder and Gibson, 2010).

Moreover, the statistics do not include failures due to application software faults. Computer centers typically keep statistics only for the failures they see themselves responsible for. With application software failures included, the fraction of failures due to software faults is likely to be much higher.

Unfortunately, failures due to software are less well tracked and characterized. While statistics may indicate which subsystem crashed (e.g. file system), they do not indicate why the file system crashed. Therefore, much of the discussion in this section is qualitative.

4.1 Classes of software faults

Software faults can be grouped into three categories: pure software problems, hardware problems mishandled by software, and software causing a hardware problem.

4.1.1 Class 1: Pure software errors. Some of the software faults in the first category are ‘classical’ correctness issues: unhandled exceptions, incorrect return values, including null objects, and incorrect control flows, such as some function not being called or called under a different condition from what was desired. Such errors are likely to be frequent in the exascale system software stack. It is well known that system software is harder to develop than application software, kernel software is harder to debug than user software, and *reactive* software, where execution is driven by asynchronous events, is harder to get right than is *transformational* software, such as scientific software, that transforms an input into an output through a long sequence of (mostly) deterministic transformations.

Large scale is worsening the frequency or impact of two other types of software error: concurrency and performance.

Concurrency errors: Subsystems such as a parallel file system are large, concurrent applications. Concurrent code is hard to develop because programmers have difficulty comprehending the possible interactions between a large number of agents. Humans often are said to be able to conceive of concurrency only at a limited scale (roughly up to 10), much less than the scale of large supercomputers. Concurrent code also is hard to debug because of the large number of possible interleavings of actions. Debugging tools typically are designed to handle bugs caused by the interaction of only two or a few agents. Because of the large number of agents in supercomputers and their tight interaction, failures due to subtle interactions between many agents become more frequent. The problem is compounded by stringent performance requirements that prevent the use of simple, coarse-grained synchronization. As an example, early versions of the Luster file system would occasionally corrupt the data written on files (Hedges et al., 2005).

Performance errors: By ‘performance errors’ we mean failures due to resource (time, memory, etc.) exhaustion. These manifest themselves in the form of

unacceptable performance, or actual crashes, due to timeouts (“time overflow”) or buffer overflows. Current programming models and programming methodologies do not provide good ways to manage the performance of large, distributed systems. Estimating the average load on different nodes is relatively easy, but understanding the tail of a distribution and evaluating the frequency of rare events is much harder. A large system is a ‘black swan detector’: events that occur rarely on one node are much more frequent with 1,000,000 nodes. Unfortunately, humans are not good at handling the impact of ‘black swans’ (Taleb, 2010). The Luster file system has suffered from multiple performance errors when deployed at large scale (Shipman et al., 2010). Some of the problems were due to a lack of clarity on the ‘acceptable performance behavior’ of applications. Programming models do not prevent applications from bringing a system down by taxing particular resources. In the case of Luster, one ‘Achilles heel’ was a limited ability to handle metadata operations. The designers of Luster assumed that no application would open or close tens of thousands of files each second; some applications did the unthinkable.

4.1.2 Class 2: Hardware propagating up to software and software not handling it correctly. Examples of the second category are a node failure not being handled by software at other nodes (node goes down, the reliability, availability, serviceability (RAS) system notices it, but the application does not take that into account); and a disk failure causing file system failure. These kinds of failures can be seen as software faults (bugs) because the software is supposed to overcome such hardware failures. In practice, many failures seem to be due hardware errors that were mishandled by software. One plausible reason is that testing code that handles failures is difficult. Another is that software is typically designed to handle clean, fail-stop hardware failures but will be taxed by messy, intermittent errors or other strange hardware behavior.

4.1.3 Class 3: Software creating a problem for the hardware. Incorrect firmware, for example misbehaving thermal control firmware, can damage hardware; this can be seen as a firmware fault. Software can trigger an unusual usage pattern for the hardware, causing hardware errors; this can be seen as a hardware fault. In both cases, however, the software is actually the culprit.

4.2 Severity of software faults

Not all software errors are equally bad. The syslog standard (RFC 3164 / RFC 5424) (Network Working Group, 2009) defines eight levels of severity:

- 0 Emergency: system unusable
- 1 Alert: immediate action required
- 2 Critical: critical conditions
- 3 Error: error conditions
- 4 Warning: warning conditions

- 5 Notice: normal but significant condition
- 6 Informational: informational messages
- 7 Debug: debug-level messages

Other dimensions are important as well; in particular, one must understand the scope of an error: how the error propagates and what it affects. Errors with a local effect are much easier to handle than errors that have a global effect. In software, we want as much as possible to avoid errors that corrupt the large, global system state, where recovery will involve the entire system and may take a long time.

4.3 Evolution of failure types and rates at the exascale

We expect a significant increase in software faults as we move to exascale. The software stack will become more complex as it has to handle more issues (such as power management, resilience, and heterogeneity) and has to face ever more stringent performance constraints (including memory footprint). Correctness bugs will be more numerous. The increasing scale of such systems will certainly increase the frequency of concurrency errors and of performance errors.

The problem is compounded by obstacles due to the development process for extreme-scale software. Supercomputing is a small market; the development of software for the largest systems is usually underfunded and understaffed. Furthermore, software for the largest systems is never tested at full scale before they are deployed: vendors cannot afford to stand test systems at full scale, and full-scale testing is done on-premise. As systems keep increasing in size, new software errors will surface with each new generation of systems, even if the software does not change.

5 Error prevention, detection, and recovery

In the preceding two sections we discussed sources of errors. error-handling can be categorized under several headings.

Prevention While an error-free system is not within the realm of possibility, various techniques can reduce the occurrence of errors.

Prediction Certain patterns of behavior can indicate future errors. If future errors are predicted with high precision, then preventive actions can be used to avoid them.

Tolerance Various techniques can be used to ensure that errors do not lead to failures, even if they are not detected.

Detection If an error cannot be tolerated, then it must be detected before it can be corrected.

Containment error-handling is eased if errors are contained so that they affect only a small part of the system.

Recovery Once an error is detected, forward or backward recovery is used to bring the system back to a correct state. Recovery will most often be automated.

Diagnosis As part of error detection and recovery, or at a later time, diagnosis activities can narrow down the likely cause of an error.

Repair The recurrence of errors can be avoided by replacing components, updating software, changing configuration parameters, and so on.

We address each of these approaches in the following subsections.

5.1 Prevention

We discussed in Section 3 mechanisms for hardening hardware and avoiding hardware errors. Suitable codes can be used to detect and correct errors in memory, caches, and buses. Errors in combinatorial circuits and latches can be detected and corrected by re-executing instructions.

Such prevention mechanisms can be used selectively. For example, one could have more reliable or less reliable cores, using either different designs or different operation parameters (clock speed, voltage); one could have the ability to run cores in tandem, comparing their outputs (to the L2 shared cache) in order to detect errors; one could modify mechanisms for thread-level speculation or for transactional execution so as to allow re-execution of code blocks when an error is detected; and one could have more reliable or less reliable memory. Some of these choices (e.g. types of memory or cores) need to be made when hardware is configured. Others (e.g. voltage levels, clock speeds, or duplicate execution) can be selected dynamically.

Automatic compensation mechanisms for hardware faults sometimes lead to poor overall system performance. Examples of such scenarios can be found in prior anecdotal fault analysis of large-scale systems. Sandia's Redstorm large-scale runs were plagued by slower-than-expected performance due to several CPUs running at 2.0 GHz instead of 2.2 GHz. Another Sandia system, Thunderbird, experienced poor system performance due to several InfiniBand links silently degrading to 256 MB/s instead of 1 GB/s. The tightly coupled nature of supercomputers exacerbates these issues, leading to the entire system experiencing performance loss as a result of a small set of degraded components.

One proposed approach is pervasive self-test diagnostics that run before and potentially during application execution in order to ascertain the health of system components as well as the overall system (Kerbyson et al., 2012). Similar diagnostics are run during system bring-up and in some cases weekly as part of scheduled maintenance windows, but systemic errors and performance degradation caused by transient faults

happen at a much finer granularity because of various causes, including environmental variability, certain workloads exercising components of the system in unusual ways, and human error. The more pervasive use of such diagnostics would enable a consistent performance environment from run to run, eliminating significant variability in application performance resulting from latent undiagnosed system issues.

The tradeoff here would be the overhead of running diagnostics at boot time and periodically during execution versus the possibility of performance degradation. Some of this overhead may be mitigated by a '+ 1' core whose operation will not significantly interfere with the actual workload running on the other cores. Finding the right balance between background monitoring, periodic health diagnostics, and other forms of online self-test will be an important aspect of co-design research on extreme-scale systems.

A complementary approach to software-based errors would be to adopt better design and testing methodologies. For example, performance errors could be avoided by adopting techniques used in the design of real-time software for avoiding overcommitment of resources. Alternatively, resource exhaustion could be avoided by the use of properly designed feedback mechanisms, derived from a principled application of control theory.

5.2 Prediction

A failure can be prevented by predicting the faults that cause the failure and evading the failure. For example, if one can predict that a node is likely to fail, then one can prevent job failure by vacating the node and migrating its workload to another node before the failure occurs. To do so, one needs to understand which faults are most likely to cause failures, and one needs to predict the occurrence of such faults based on past observations. The prediction should be timely: it is easy, but not very useful, to predict that each piece of hardware will eventually fail. Conversely, if the prediction is too close in time to the failure, then there may not be enough time for evading the failure. Failure prediction is used successfully for a wide range of complex systems, including railroads (Oh et al., 2006), nuclear power plants (Zio et al., 2010), and aircraft engines (Hunter, 1975). Many different techniques can be used to forecast failures. A fairly complete survey of these techniques is presented in Salfner et al. (2010).

Several studies suggest that failure can be predicted in HPC systems. For example, a memory device tends to show, for a given address, multiple repetitive correctable errors before showing an incorrectable error (Hwang et al., 2012). Correlations in time have also been observed between soft errors and hard errors. Another recent study (Heien et al., 2011) has observed correlation in space. The predictability of hard drive failure is at the

origin of the Self-Monitoring, Analysis and Reporting Technology (SMART) used in many disks.

In HPC systems, the overall failure prediction workflow based on event analysis, its limitations, and needed improvements are reasonably well identified. HPC systems are producing events related to the state of their software and hardware components. Events of the same type can be clustered into groups. Event correlation analysis allows establishing stochastic propagation chains between events of the same group and/or of different groups. Stochastic propagation chains essentially contain two categories of events: precursors and critical events. When a critical event is in a propagation chain, all previous events in the chain are called precursors (precursors potentially also include critical events). In the past two years, several key results have demonstrated that recent advances in event clustering (Gainaru et al., 2011b), anomaly detection (Gainaru et al., 2012a), event correlation (Gainaru et al., 2012b), propagation chain construction (Heien et al., 2011), and online detection of propagation chains (Gainaru et al., 2011a) can provide precise failure prediction. The time lag observed for the most efficient prediction approaches is consistent with the time taken by proactive actions.

Current predictors can achieve a precision of over 90%, so that preventive actions will be superfluous in only one-tenth of the cases; acting on such predictions is usually worthwhile. On the other hand, the recall is still low and stays below 50% even for the most advanced prediction approaches: fault prediction can effectively double the MTBI but cannot replace other methods, by itself. The main reasons for the low recall are the lack of precursor events (some failures have no identified precursors) and the precision losses at each step of the failure prediction workflow. Thus, an identified research objective is to improve the whole failure prediction workflow to increase the failure prediction coverage from 50% to 80% or 90%.

5.3 Tolerance

For some applications, we may not need to recover from node failures at all. For example, in derivative-free parameter estimation of a complex simulation, a node failure could be ignored and treated as a simulation failure. However, not all simulation failures are the same. A graceful failure can yield partial information that could be used when determining the next experiment to perform for the optimization. Structured simulation-based optimization techniques can use this partial information to build partial interpolation models and thus become resilient to node failures. Similarly, we could use partial solutions for simulations at a looser tolerance as long as we account for the truncation error in the model and optimization.

This approach can be likened to controlling the noise in simulations (Moré and Wild, 2012). For stochastic

noise, model-based optimization methods have been developed that specify both a candidate point and the number of replications needed to obtain sufficient accuracy. Parallel replications at a fixed point can be used to control stochastic noise but not deterministic noise. For deterministic functions one could use nearby points and Taylor's theorem to bound the noise in the simulation. By neighborhood-sampling one could reduce the noise in many settings, and these samples may already be available from previous computations of the algorithm.

An alternative approach is to use insight into the application to reduce the probability of failure by using a smaller word size. Variable precision arithmetic can help in this approach by using bounds on the precision requirements for Newton solves to compute low-precision steps initially. Analysis tools, such as those developed for automatic differentiation and estimating computational noise, could identify blocks in the code for which higher precision would lead to improved precision in function evaluations. Based on this identification, one could restructure the computation of a function so that the least-precision arithmetic was used in each block to obtain the required precision in the overall function evaluation. Similar ideas could be applied for gradient and Hessian evaluations. User-provided and automatically generated codes for quantities derived from function values (such as derivatives) can be significantly less precise than the underlying function. Analysis of the underlying computational graph (for example as done by (Kubota 1992)) could provide insight into reformulations of the derived code that yield both function and derived values to specified precision. The use of a smaller word size reduces the number of gates and latches involved in computations, thus reducing the frequency of errors.

5.4 Detection

Mechanisms for detecting hardware errors, such as ECC and circuit-level redundancy, are briefly described in Section 3. Here we focus on software-driven detection and application-level detection.

5.4.1 Software-driven detection of hardware errors.

Conventional hardware detectors either have relied on expensive redundancy-based solutions or have focused on specific fault models and hardware components. Recently, considerable work has been done on software-driven solutions that are oblivious to the fault model and potentially provide larger hardware coverage at low cost. The key observation underlying these techniques is that the hardware reliability solution needs to handle only those faults that become observable to software. This class of solutions, therefore, focuses on detecting hardware faults by monitoring for anomalous software behavior or symptoms. Much

research has shown that such monitors (implemented in software and/or hardware) can be inexpensive and detect a wide range of hardware faults (Goloubeva et al., 2003; Pattabiraman et al., 2006; Wang and Patel, 2006; Dimitrov and Zhou, 2007; Meixner et al., 2007; Racunas et al., 2007; Li et al., 2008b; Hari et al., 2009; Lyle et al., 2009). Moreover, this strategy treats hardware faults as analogous to software bugs, potentially leveraging software reliability techniques and further amortizing overheads.

A software-anomaly- or symptom-based detection strategy must be viewed in the context of a holistic reliability solution. First, since the hardware fault is detected through software symptoms, the latency from the activation of the fault to detection can be high (relative to traditional hardware-driven techniques). This requires a sophisticated diagnosis strategy to determine the root cause of the symptom, namely whether it was a hardware or a software fault; in the case of a hardware fault, whether it was a permanent or a transient fault; and in the case of a permanent fault, in which field replaceable unit the fault occurred so as to trigger appropriate repair/reconfiguration and recovery. Simplifying detection in exchange for a more complex diagnosis is a reasonable tradeoff since the former is 'always on', whereas the latter is invoked in the relatively infrequent case of a fault detection.

Second, the longer detection latency also impacts recovery. Software-driven detection techniques rely on backward error recovery, typically checkpoint/rollback-based recovery. Therefore, the detection latency should be short enough to ensure that a fault-free (recoverable) checkpoint is available on detection. Another constraint comes from the need to buffer outputs until they are known to be fault-free; the detection latency should be short enough to ensure that this buffering time does not degrade performance.

Much recent work has been done on individual components of the above approach (Prvulovic et al., 2002; Sorin et al., 2002; Goloubeva et al., 2003; Nakano et al., 2006; Pattabiraman et al., 2006; Wang and Patel, 2006; Bower et al., 2007; Dimitrov and Zhou, 2007; Meixner et al., 2007; Racunas et al., 2007; Li et al., 2008a,b; Sahoo et al., 2008; Hari et al., 2009; Lyle et al., 2009; Ramachandran, 2011). Recent work on the SWAT (SoftWare Anomaly Treatment) project (Li et al., 2008a,b; Sahoo et al., 2008; Hari et al., 2009; Ramachandran, 2011) has developed an integrated framework for all components of such a resiliency solution with promising results. It performs software anomaly detections using both hardware monitors (e.g. fatal traps that require no added cost or more explicit hardware out-of-bounds detectors that detect addressing anomalies) and software monitors (e.g. the kernel panic routine that involves zero cost or more explicit application-level invariant checkers). The detectors

invoke a thin firmware layer that diagnoses the root cause of the symptom, leveraging the rollback/replay mechanism available for recovery. Repeated replays on different cores and units are used to systematically narrow down the source of the fault. Once the root cause is understood and eliminated or repaired, recovery is invoked, and application execution continues.

The software-driven approach described has several advantages: (1) *generality*: it is oblivious to specific failure modes and microarchitectural or circuit details; (2) *masked faults ignored*: it naturally ignores all faults masked at the software level; (3) *customizability*: the software layer in charge of resilience can be customized to the application and system in various ways; and (4) *amortization of overheads*: the approach is inspired by online software bug detection (Hangal and Lam, 2002; Ernst et al., 2007) and can leverage similar techniques, thereby amortizing overheads toward a holistic view of system reliability.

A key limitation of the approach is that some faults could corrupt application state in undesirable ways but escape detection. Such SDCs could potentially be catastrophic, and much research is required to mitigate their effects. A key problem is that the conventional method to quantify the presence or impact of SDCs relies on fault (or error) injection campaigns (using real applications; see for example Reis et al., 2005a; Wang and Patel, 2006; Li et al., 2008b; Lyle et al., 2009). With the above approach, the impact of a fault depends on the application and where in the application the fault was injected. A brute-force fault injection campaign might require trillions of fault injections (one fault per application and hardware fault site) even for simple benchmarks and hardware fault models and is clearly impractical. Therefore, statistical fault injection campaigns are used where a random sample of application instructions (and hardware sites) is selected for fault injection, but these do not provide any insight on where (if) SDCs might occur in the rest of the application. Without such knowledge, it is difficult to design protection mechanisms for the SDC-vulnerable parts of the application.

Significant progress has been made recently in addressing this problem. For example, recent work on Relyzer (Hari et al., 2012a,b) proposes methods to determine when application-level transient faults are equivalent, enabling comprehensive analysis by injecting (transient) faults in only one instruction per equivalence class. Relyzer is able to both determine all SDC-vulnerable fault sites with relatively high accuracy for the studied fault model and identify the reason for the SDC (i.e. the fault propagation path). The latter motivates low-cost, application-specific detectors designed to protect only those instructions that are vulnerable, thereby enabling selective, frugal, and customizable placement of detectors. The approach promises quantifiable resiliency vs overhead tradeoff

curves that can be used as appropriate by the system designer or application writer. Another project, SymPLFIED (Pattabiraman et al., 2008), takes a complementary approach of understanding the impact of different errors in the same application site without performing different fault injections for each. SymPLFIED inserts a symbolic error value and uses model-checking to explore all execution paths with this value, ensuring that all paths that result in corruptions are detected and, if not, to motivate detectors. This approach has been tried only for relatively small programs, however, because model-checking is resource-intensive. The Shoestring project (Feng et al., 2010) has developed a pure static analysis that identifies instructions where faults are likely to be detected quickly enough (e.g. there is a short path in the data-flow graph from such a fault to enough potentially symptom-generating instructions) without requiring fault injections. The rest of the faults are considered vulnerable and protected by using selective instruction duplication. Shoestring reduces the SDC rate significantly but is not yet able to eliminate SDCs or comprehensively identify where the remaining ones are.

Despite this progress, much research remains to be done to convert ideas such as these into a practical workflow that can be demonstrated for all fault models of interest and that can drive automatic derivation and insertion of detectors according to customizable resiliency vs overhead tradeoff requirements.

5.4.2 Application-level detection of hardware errors. At the application-software level, we can develop a taxonomy of errors similar to the one presented in Section 2. We separate errors into detectable and undetectable errors. An example of an undetectable error is a corrupted matrix/vector dimension before we invoke a checksum. We can further subdivide each category into irrelevant errors (such as errors in out-of-date data that will not be used further), correctable errors (such as a single corrupted matrix element that can be corrected using checksum), and uncorrectable errors. The key message is that although application-level detection can handle some hardware errors, it cannot, on its own, ensure resiliency. At the same time, application-level detection can mitigate the overhead of error correction in hardware or lower-level system software and thus forms part of an integrated approach to resiliency.

Application-level error-detection schemes have been developed in the context of solvers for linear systems (Huang and Abraham, 1984; Banerjee and Abraham, 1986; Banerjee et al., 1990; Turmon et al., 2000; Gunnels et al., 2001; Turmon et al., 2003; Chen and Dongarra, 2006; Bosilca et al., 2009) and certain iterative methods for solving partial differential equations (PDEs) (Roy-Chowdhury et al., 1996). These schemes are based on computing checksums of the rows and/or

columns of the matrix (discretized PDE). The checksums can be shown to preserve a range of common matrix operations such as addition, multiplication, scalar product, and LU and QR decomposition (i.e. matrix inversion or solves). Checksums can thus detect errors in common matrix operations, although strictly speaking we can detect only the fact that the checksum is inconsistent, which may indicate a corrupted matrix element or an error during the checksum or matrix operation (a common misconception). With this caveat, a single erroneous matrix element can be corrected by using checksums (more elements can be corrected if the matrix decomposes and the errors occur in independent partitions). Unfortunately, the checksums have not been generalized to multigrid methods (Hackbusch, 1985; Trottenberg et al., 2001) for solving PDEs, which are optimal in terms of flop counts compared with the SOR method described in Roy-Chowdhury et al. (1996). Also, some care is needed to define tests that ignore normal round-off errors but catch most silent hardware errors (Turmon et al., 2003).

Application-level detection in other areas of applied mathematics is less well developed (in part, this situation may be because other areas such as optimization or differential equations can be built on resilient linear algebra routines, provided the remaining computations are performed in a resilient manner). However, additional opportunities exist at higher levels of abstraction to design resilient algorithms at potentially reduced overheads. For example, when we are solving a nonlinear system of equations, $F(x) = 0$, with Newton's method, we typically promote convergence by enforcing descent in a merit function such as $p(x) = \|F(x)\|_2^2$ for the Newton step, s_k , at iteration k , obtained by solving the linear system $\nabla F(x_k)s_k = -F(x_k)$. Solvers assess progress by ensuring a sufficient reduction condition in the merit function (e.g. Fletcher, 1981) such as

$$p(x_k) - p(x_k + s_k) \geq \sigma (\|F(x_k)\|_2^2 - \|\nabla F(x_k)s_k + F(x_k)\|_2^2)$$

where $\sigma \in (0, 1)$. We can use this condition to detect errors during the computation of the (approximate) Newton step. If the right-hand side is negative, then the solve failed. If the sufficient reduction condition fails, then we recompute the Newton step inside a reduced trust-region (e.g. Conn et al., 1987).

Other application-level error-detection schemes can be built around invariants. For example, we can detect errors in the gradient computation, $\nabla F(x_k)$, by recomputing gradients of $p(x_k)$ at a cost that is comparable to a single function evaluation using automatic differentiation (Griewank and Corliss, 1991) to detect errors ($\nabla p(x_k) = \nabla F(x_k)F(x_k)$). Stochastic optimization (Birge and Louveaux, 1997) and stochastic PDEs (Chow, 2007) also provide error-detection schemes. In both

cases, we typically solve an ensemble of systems and compute expected values. Thus, we can use the deviation from the expected value to detect potential errors in individual ensembles. However, such a scheme cannot detect all errors (e.g. those that are close to the expected value). An interesting challenge is the integration of hardware error models into the convergence analysis of these sampling methods.

Invariants can also be derived from the physics of the simulated system: wind speed is positive and does not exceed the speed of sound; nearby values cannot be too different; system energy is preserved. Programmers often check such invariants in order to debug their codes; using such checks to catch hardware errors may not add much coding effort.

The two approaches described in the preceding two sections (software-driven detection and application-level detection) are nicely complementary. Software-driven detection is most effective for errors that affect the system state or the control state of an application (e.g. wrong jump) or break the language abstractions (e.g. corrupt pointers); application state detection is most applicable when the application computation proceeds unperturbed, but data values are incorrect. Research is needed to further study this complementarity and understand the coverage obtained when both methods are used.

5.4.3 Behavioral-based detection. The number of cores used in large-scale systems already exceeds a million cores. As a result, the challenge of developing correct, high-performance applications is also growing. When an application does not complete or completes with incorrect results, the developer must identify the offending task (such as an MPI task) and then the portion of the code in that task that caused the error. Traditional parallel debugging tools (Lourenço and Cunha, 2001; Lindekugel et al., 2008; MPIPlugIn, 2013; Rogue Wave Software, 2013) often perform poorly at large task counts. Hence, research is actively underway to develop a detection toolchain that can identify the offending task and, to a customizable granularity, the relevant portion of code within the task responsible for the error.

Several debugging tools detect bugs in large-scale applications without relying on extensive manual effort demand by debuggers such as gdb, DDT, or TotalView. These more sophisticated debugging tools typically focus on detecting violations of deterministic and statistical properties of the applications. Deterministic tools can validate certain properties at runtime; any violation of these properties during an execution is reported as an anomaly. For example, FlowChecker (Gao et al., 2010) focuses on communication-related bugs in MPI libraries. It extracts information on the application's intentions of message passing (e.g. by matching MPI Sends with

MPI Receives) and at runtime checks whether the data movement conforms to these intentions. Bug localization follows directly: the data movement function that caused a discrepancy is the location of the bug.

Statistical tools (Mirgorodskiy et al., 2006; Gao et al., 2007) detect bugs by deriving the application's normal behavior and looking for deviations from it. For example, if the behavior of a process is similar to the aggregate behavior of a large number of other processes, then it is considered correct, and different behaviors are considered incorrect. Mirgorodskiy et al. (2006) monitor the application's timing behaviors and focus the developer on tasks and code regions that exhibit unusual behaviors. This approach centers on function call traces in order to identify the trace that is most different from other traces. DMTracker (Gao et al., 2007) uses data-movement-related invariants, tracking the frequency of data movement and the chain of processes through which data moves.

While these tools are effective in their own domains, their primary weakness is that their designs do not consider scalability. Typically, these tools collect trace data during the application's execution and write it to a central location. They then process the data in order to detect potential problems. Some recent work has tried to rectify this problem by analyzing the application's behavior online, without any central bottlenecks. One such work is STAT (Lee et al., 2007, 2008; Ahn et al., 2009), which provides scalable detection of task equivalence classes based on the functions that the processes execute. STAT uses MRNet (Roth et al., 2003), a tree-based overlay network, to gather and merge stack traces across tasks and presents the traces in a call-graph prefix tree that identifies task equivalence classes. STAT removes problems associated with a central bottleneck by reducing the trace data as part of a computation being performed within the overlay network through a custom reduction plug-in.

Another work of this type is AutomaDeD (Bronevetsky et al., 2010; Laguna et al., 2011, 2012), which performs runtime monitoring of a parallel application to build a statistical model of the application's typical timing and control-flow behavior. AutomaDeD models the control flow and timing behavior of application tasks as semi-Markov models (SMMs) and detects faults that affect these behaviors. AutomaDeD examines how each task's SMM changes over time and relates to the SMMs of other tasks in order to identify the task and code region where a given fault is first manifested. AutomaDeD detects which time period in the execution of the application is likely erroneous. Next, it clusters task SMMs of that period and performs cluster isolation, which uses a novel similarity measure to identify the task(s) suffering from the fault. Then, transition isolation detects the transitions that were affected by the fault more strongly or earlier than

others, thus identifying the code region where the fault is first manifested. STAT focuses primarily on the state of the application once an error manifests itself, whereas AutomaDeD focuses on scalable analysis of the entire application execution.

While behavioral-based detection work has been focused on debugging, the identification of anomalous behaviors can be used to detect errors that are due to other causes as well. Further work is needed to make the behavior-based detection tools robust enough to rely on in-production systems. One question is how a detection system should deal with changing workload patterns, and corresponding discontinuous, but legitimate, changes in the correlation patterns. A related question is how a detection system should handle noise in the execution environment, such as that resulting from congestion on the network switches due to competing applications executing on other nodes. For purposes of scalability, tools compress models for comparison. It is tempting to use lossy compression for this purpose. If so, what parts of the model can be compressed away because they are not germane to the error detection or localization activities? Moreover, are the models powerful enough to handle a wide variety of applications and their legitimate behaviors and yet simple enough that their parameters can be reliably derived through training and detection, and localization can be done efficiently at runtime using the models?

5.5 Containment

Most system-level failures affect only a single node. Statistical analysis (Bautista-Gomez et al., 2011a) shows that multinode failures, also called correlated failures, are rare. The probability distribution of multinode failures according to the number of nodes involved in a correlated failure is heavy-tailed: failures involving the whole system are rare but still happen, for example in the case of a long power outage.

On the other hand, the global checkpoint/restart approach to application recovery makes the simplifying assumption that if an error occurred, then any application state could be corrupted. In practice, by the time an error is detected, it may have propagated to only a small subset of the application state. Recovery could be faster if only this small fraction of the application data was repaired.

5.5.1 Strategies to limit propagation. Various containment strategies can be used to limit error propagation.

A priori containment recursively divides the resources of a parallel system and execution of a parallel program into nested disjoint containment domains (CDs); the goal is to limit recovery to one CD, at the finest nesting granularity possible. Any error or failure

can be contained within some level of the CD tree and may be recovered by restoring only the state necessary to re-execute that CD. State is restored from explicit preservation clauses within each CD, which permit a range of preservation/restoration tradeoffs. These include preserving only a partial state, relying on regeneration routines or on state already available elsewhere in the system or at a higher CD level. Alternatively, one could use forward recovery of state where the state of a CD is corrected, for example by extrapolating from the state of neighbors; this is discussed in Section 5.6.4. These approaches can be applied hierarchically. If recovery fails at one level of the system it falls back to recovery at a higher level (which, presumably, is more expensive but more reliable) (Chung et al., 2012).

The choice of CDs in terms of granularity, preservation/restoration options, and recovery and detection routines introduces new, flexible tradeoffs. For example, one can construct a strict CD hierarchy in which all communication occurs at a single CD context at a time, simplifying preservation and recovery. Often, however, it is preferable to relax this communication constraint in order to reduce preservation overheads and the granularity of recovery. When communication is allowed between CDs, data must first be verified for correctness in order to prevent silent data correction. Communication should also be logged in order to retain the ability to recover CDs in an uncoordinated manner (see Section 5.6.2). Overall the tradeoffs are between the cost of preserving state (lower relative overhead for larger domains) and the cost of CD recovery (which is relatively higher if containers are large).

CDs can be selected statically, based on the application structure and tuned automatically for optimal resilience. For example, in a multiphysics code, modules running the different physics codes are natural CDs, with the containment done by the coupler that couples these modules. Alternatively, one can build CDs automatically by tracking communication during a trial run and finding good separators in the communication graph (Ropars et al., 2011).

Another possible approach is a posteriori containment. The logic of the application may constrain error propagation. For example, in an iterative algorithm with nearest-neighbor communication, an error can propagate at most one neighbor away at each iteration. In a 3D problem, where each node holds a $k \times k \times k$ subcube, a bit flip in a data element will have propagated to at most $(n/k)^3$ nodes after n iterations. An algorithm that checks periodically for corrupted values can compute a posteriori the domain that could be affected by their error and use localized recovery.

Another form of a posteriori containment is the retention of multiple checkpoints (Hogan et al., 2012; Lu et al., 2013) and recovery based on analysis or more extensive error-checking than one would normally

incur at each checkpoint. When an error is manifested and the system proceeds to recovery, then based on application semantics, the contents of multiple checkpoints can be analyzed to find the most recent one that has a correct application state that can be resumed. This is a particularly powerful technique for tolerating silent errors and may avoid the need for checking checkpoints for correctness as they are committed. Such multiversion checkpoints are likely to be most viable for applications that have modest main memory requirements, when application-level determination of critical state is employed, or when additional resources such as NVRAM are available.

5.5.2 System software. Application recovery requires a correct functioning of multiple global system services (resource managers, parallel file system, etc.). Failures in these systems are much harder to recover from and often require a time-consuming reboot. Thus, containment techniques are important for OS functions.

One approach to this problem is to partition system software in such a way that even when corruption occurs in systems code, it can be contained, and failures in a particular core do not impact other cores. Such a partitioning can be accomplished within an OS image by taking a formally verified microkernel approach with the system software (Heiser et al., 2011), by using a hypervisor such as Palacios (Lange et al., 2010) or a hybrid kernel approach such as those proposed by NIX (Ballesteros et al., 2012) or FusedOS (Park et al., 2012). All of these approaches create strict boundaries between different software components of the system, which can be used to facilitate the creation of CDs within the system services and applications. Selective restart or fail-over of those partitions can refine the granularity of recovery to improve efficiency.

5.6 Recovery

Recovery will return the system to a valid state. Backward recovery returns the system to a previous state (a previous checkpoint), whereas forward recovery evolves the system to a new, correct state. Currently, in high-performance computers, system state is recovered by forward recovery, while application state is recovered by backward recovery. Checkpoint/restart is advantageous when large parts of the computation state change rapidly; this is the case with application variables in a scientific computation. Replication at MPI process level has been explored (Ferreira et al., 2011). Its cost is high, and this approach is competitive against checkpoint/restart only in extreme situations. Full-node replication has not been explored in the HPC domain as far as we are aware. Its cost in development and overheads would be even more expensive than replication at the level of MPI processes.

Forward recovery makes sense when a relatively small part of the state changes; this is the case with a file system and with the system state (most of which does not change during a scientific computation). Forward recovery requires sufficient redundancy in stored state that a correct state can be recreated if part of it was lost. It also requires the use of update mechanisms that ensure that a failure in the midst of an update will not corrupt the state. Commit protocols and transaction logging for replay are two examples.

Research efforts in this area focus on avoiding the need for a global checkpoint/restart for applications, by ensuring that errors are contained and recovery can be performed locally. If the OS and runtime have a more dynamic behavior (e.g. resources added or deleted during a computation, processes migrated), then forward recovery of the OS and runtime will require additional effort.

5.6.1 Restart. The classical checkpoint/restart strategy for resilience used in most large-scale executions in petascale systems has two main limitations: (1) the time to save the state of the execution (checkpoint) is becoming unacceptable compared with the system MTBF, and (2) all processes involved in the execution are restarted from the last checkpoint even if only one process fails. Recent results in multilevel checkpointing and in fault-tolerance protocol show that these two limiting factors could be addressed and make checkpoint/restart a viable approach for exascale resilience for errors that are quickly detected (detected in less time than it takes to commit a checkpoint).

Multilevel checkpointing (hybrid checkpointing) consists of using multiple storage resources with different characteristics in terms of speed and reliability in order to respond to different failure scenarios. The main scenarios to consider are the crash of a process that can be restarted on the same node, the failure of a node that makes that node unavailable for restart, and the failure of the entire system.

Multilevel checkpoint restart uses local storage resources (NVM, hard disk drive (HDD), or solid-state drive (SSD) devices) as a first level of storage for execution checkpoints. A second level could use the storage resources of remote nodes. If a node failed, even if it cannot be restarted, the execution context of that node could be restarted from the checkpoint stored on remote node. Local, persistent storage can also handle node failures if it is twin-tailed, that is, remotely accessible even after a node failure. A third level of checkpoint considers an encoding of several process checkpoints and a distributed storage of the encoding result on several nodes. Different encoding algorithms (Xor, Reed Solomon, etc.) can be used, leading to different levels of reliability. According to the level of reliability provided by the encoding algorithm, this third level of checkpointing can be used to

tolerate simultaneous multinode failures. A fourth level of storage is the remote parallel file system. This level is relevant only for catastrophic failure scenarios that could not be covered by the previous checkpointing storage levels, such as the loss of enough nodes to make the restoration of the checkpoint images impossible. Finally, mass storage can back up disk information, enabling recovery from catastrophic failures of the file system. The available bandwidth for checkpoint storage of several levels of storage is studied in Moody et al. (2010).

Currently, two environments provide multilevel checkpoint/restart: SCR (scalable checkpoint/restart) (Moody et al., 2010) and FTI (fault-tolerance interface) (Bautista-Gomez et al., 2011b). While SCR is keeping the file interface abstraction, FTI is providing a data structure abstraction, masking from the programmer how the data to be saved is actually managed by the library. Recent results show that a process context of 1 GB can be saved in 2–3 s in local SSD (two SSDs mounted in RAID0). Such checkpoint speed is orders of magnitude faster than checkpointing on a remote file system, which requires tens of minutes on current petascale systems (about 30 mins if the full system memory is dumped in the remote file system). An experiment with FTI on a current large-scale execution (0.5 million GPU cores) of an earthquake simulation on a hybrid system composed of CPUs and GPUs demonstrates very low overheads on the execution time (less than 10%) when using a checkpoint strategy, compared with a computation that does not checkpoint. Other research results demonstrate that checkpointing on remote node memory is even faster than on local HDD or SSD (Zheng et al., 2012). Research still is needed, however, in order to understand how to take advantage of new storage technologies such as phase change memory. Europe has a project called advanced multilevel fault tolerance (AMFT) to test this approach with several storage technologies; the objective is to include multilevel checkpoint restart in the PRACE software stack and to prepare for exascale.

5.6.2 Localized restart. Checkpoint/restart is usually done at the application level. Applications periodically save state onto storage and provide a callback function to restore the computation from saved state. For most applications, the checkpoint size is a fraction of the system memory. Checkpointing is coordinated: the involved processes synchronize before checkpointing and ensure that no message is in flight during the checkpoint operation.

If the computation is restarted, then all processes restart from the last checkpoint, even if only one process has failed. In general, this situation cannot be avoided. If the computation is nondeterministic, the computation after restart could follow a different path from that followed before the failure occurred; the

computations of the ‘healthy’ processes may not be valid anymore. However, most HPC scientific codes are ‘piecewise deterministic’: the execution consists of long deterministic phases, with nondeterminism occurring at a small (possibly empty) set of execution points. Thus, the opportunity exists to use message-logging protocols in order to avoid global restarts. During the fault-free execution, all messages contents and nondeterministic events (reception orders) are recorded. When a failure occurs, only the failed process restarts; its state is reconstructed by sending it the messages recorded before the failure and by forcing the message deliveries in the same order. Many variants of message-logging protocols have been developed (Elnozahy et al., 2002). However, they all share two limitations: (1) the contents of all the messages need to be saved, requiring a significant amount of storage; and (2) the nondeterministic events (reception orders) also need to be stored, thus impacting either the communication latency or the communication bandwidth, depending on the message-logging protocol.

A recent analysis of communications patterns in HPC applications shows two important properties: (1) communication patterns are either deterministic (the order and outcome of communication operations at each process are fixed) or send-deterministic (whatever the order of reception for each process, the sequence of send operations is identical in any correct execution) (Cappello et al., 2010); and (2) communications show strong spatial and temporal localities and form clusters, which can be observed manually for some applications and extracted automatically with graph-partitioning tools (Ropars et al., 2011). These two properties can be leveraged to develop new fault-tolerant protocols having excellent properties in the HPC context: no global restart, no need to log all message content, no need to store reception orders, no risk of restart from the beginning. Two fault-tolerant protocols have been proposed in the literature (Guermouche et al., 2011, 2012) from these principles. A hybrid protocol can use coordinated checkpointing inside clusters and message-logging between clusters. This protocol is a good match for HPC applications built of independent modules, such as the CESM climate simulation code (NCAR, 2014): checkpoint/restart can be done independently for each module (cluster), and logging (within the coupling toolkit) handles interaction within modules. For real applications, the number of messages to log is a small fraction (10%) of all the messages sent during the execution (Guermouche et al., 2012). Other hierarchical, hybrid fault-tolerant protocols, combining coordinated checkpointing with some form of message-logging, have been proposed that do not consider communication determinism (Bouteiller et al., 2011). They require logging, in some way, the message reception orders of all messages.

While this progress is encouraging, many research questions remain open: how to form clusters to reduce the number of messages to log, how to adapt clusters to the different communication patterns seen during the execution, how to prove the deterministic or send-deterministic nature of communication patterns automatically, how to organize a fully distributed recovery, how to better understand sources of nondeterminism in applications that show nondeterministic or send-deterministic communication patterns, and how to address them.

Localized restart reduces the total I/O volume needed to restart, but it may not reduce the restart time if all nodes have to wait for the failed node to be restarted. Nevertheless, it still may result in lowered power consumption, since the waiting nodes can reduce their power intake. Furthermore, the restart can be accelerated by being distributed across multiple nodes.

5.6.3 Fault-tolerant data structures. Between application-level and restart schemes, there are runtime-level techniques for redundancy and repair. These techniques can operate at the data structure level, below even a typical application abstraction, and by encoding additional information into the data structures enable them to be reconstructed in case of error. Common examples include i-nodes in filesystems, redundant virtual-physical mapping information in OS page tables, trees and lists with multiply linked structures, and redundancy-encoded arrays and data structures. These techniques offer the potential for significant recovery capability under software (compiler, runtime, OS, even application) control, and they support selective and flexible usage. One example of such structures has been proposed in the Global View Resilience (GVR) system (Fujita et al., 2013).

5.6.4 Application and algorithmic recovery. Application and recovery techniques can use the algorithmic redundancy available in many parallel algorithms, in order to recreate a valid computation state if the loss of a (small) part of the state has been detected. Many simulations use iterative methods on meshes. When a catastrophic node failure occurs and is communicated to the remaining nodes, such a method could approximate the missing information and continue with the computation, by extrapolating the missing information from the remaining information. If the method had suitable convergence properties, then the error thus introduced would be smoothed out, possibly at the price of additional iterations. More sophisticated recovery methods that use a hierarchy of meshes generated for multigrid methods could also be developed. These methods would traverse from fine to coarse and back using restriction and interpolation operations. By moving to a coarser level,

one could estimate the numerical values of the computational node that failed, using the interpolation operation and neighboring values, and then construct a new mesh for the missing patch and apply interpolation operations. Such an approach requires knowledge of how to remesh and recover the mesh hierarchy, and possibly a rebalancing of the computations to prevent neighboring nodes from becoming a computational bottleneck.

When algorithmic redundancy is not available in the original problem formulation, it may be possible to add it with little increase in storage and computation. For example, dense linear algebra computations can be protected by adding checksums to the original matrix that are updated during computation, and periodically checked. In addition, one can checkpoint immutable outputs as they are produced (Du et al., 2012).

In applications with suitable patterns, the recovery mechanism might be less intrusive. For example, in branch-and-bound methods for mixed-integer optimization, which recursively subdivide the domain and solve optimization problems on each subdomain, a tree structure maintains the current state. As long as the tree structure was available, if a solve on a subdomain did not complete because of node failure, that subdomain could be recovered from the tree and the optimization problem solved on a different node. The same approach applies to any functional execution model, where variables are not mutable: if the evaluation of a function fails to complete, it can just be recomputed, assuming the inputs were preserved. This approach is heavily used by Hadoop to provide resilience (Dean and Ghemawat, 2008). Also, an input or an intermediate value may affect only some of the outputs, in which case it may be possible to restart only some of the computations, even if no intermediate value was preserved (Gunnels et al., 2001).

This discussion also suggests that algorithm-level checkpointing can be more efficient than system-level checkpointing. For example, to ensure that function, gradient, and Hessian computations are correct, one needs only to checkpoint the computational graph of the nonlinear functions, which is orders of magnitude less information than the values and sparsity patterns. Similarly, branch-and-bound schemes need only checkpoint the root node of each distributed solve, which can be stored by using two binary vectors.

Another advantage of algorithm-based recovery is that it may not be necessary to replay the MPI messages since the last checkpoint. For example, if a node fails during a distributed solve of $F(x) = 0$, we can simply resume the Newton iterations from the checkpoint, because we are not interested in the sequence of iterates, x_k , but the final solution. This opens the door to new hybrid methods that combine Newton with Gauss-Jacobi steps. The analysis of such methods remains an

open problem. In the context of stochastic optimization, asynchronous techniques already exist that can accommodate missing subproblem solves (Linderoth and Wright, 2003).

5.6.5 Fault-tolerant MPI. Application-level recovery needs to be preceded by system-level recovery. For example, if a node has failed, then the application has to be made aware of the failure and has to continue its execution in a well-defined environment, in order to execute the recovery code. This problem is usually addressed in the context of MPI. Can we ensure that the failure of one node will not cause processes running on other nodes to crash? How can we inform the other processes of the failure? What is the state of MPI after the crash? Projects aimed at providing a fault-tolerant MPI have been going on for over a decade (Fagg and Dongarra, 2000; Bouteiller et al., 2006), and several prototype implementations of fault-tolerant versions of MPI exist. The MPI forum has discussed several proposals for standardizing fault-tolerant MPI (Bland et al., 2012) but has not agreed yet on a standard.

A similar problem occurs for any library that requires consistent state across multiple nodes. These include mathematical libraries and I/O libraries. For each of these, we need to define what the state of the system is after a failure and how the information about the failure is propagated.

5.6.6 Rejuvenation. Software rejuvenation is meant to mitigate the problem of software aging, in which the state of the software system degrades with time (Castelli et al., 2001). The primary causes of this degradation are the exhaustion of OS resources (such as file handles or network sockets), data corruption, and numerical error accumulation. Eventually, software aging leads to performance degradation or correctness problems such as hang or crash failure. Some typical causes of this degradation are memory bloating and leaking, unterminated threads, unreleased file-locks, data corruption, storage-space fragmentation, and accumulation of round-off errors. These causes also affect HPC applications, and hence software rejuvenation is a relevant technique in our tool chest. In particular, accumulation of round-off errors is a problem in some numerical computations that appear in HPC applications (Hamming, 1987).

Software rejuvenation essentially involves occasionally terminating an application or a system, cleaning its internal state, and restarting it. This process removes the accumulated errors and frees up OS resources, thus preventing in a proactive manner an unplanned and potentially expensive system outage due to the software aging. Much research has been done in order to determine optimal times to do software rejuvenation (Avritzer et al., 2006; Grottko and Trivedi, 2007); for

example, when the load on the system is low, the amount of corrupted state is likely to be small, or a failure is impending. With an appropriate choice, the cost of system downtime can be reduced significantly compared with reactive recovery from failure.

Surprisingly, software rejuvenation has not been widely used in HPC applications. In Naksinehaboon et al. (2010), the authors argue that rejuvenation should be tried in HPC applications only at the level of individual OS kernel, rather than the entire system. They propose three scheduling strategies for rejuvenation: using the MTTF, the median of TTFs, and the reliability of the system. Based on failure data extracted from System 20 at Los Alamos National Lab (2006), they evaluate the hypothesis that rejuvenation together with checkpoint/restart can reduce the lost computation, over simply checkpoint/restart. The verdict is mixed. Only by a careful estimate of TTFs can rejuvenation give benefits. Not surprisingly, more rejuvenations quickly reach a point where they hurt overall performance.

Nevertheless, it seems worthwhile to further explore the application of rejuvenation in HPC applications. The first issue that needs to be considered is what state should be saved and ‘rejuvenated’. Related to this is how that state should be compartmentalized so that a quick rejuvenation is possible. The second issue is when to trigger the rejuvenation. In addition to the factors that have already been explored in non-HPC domains, here one must also consider the interactions of the node being rejuvenated with all the other nodes in the cluster on which the application is running. Done right, software rejuvenation holds the promise of extending the MTBF and reducing the frequency of checkpoint/restart.

6 System view of resilience

We discussed in the preceding section mechanisms for detecting hardware errors at the system software or application level. A similar interplay between the various system layers applies to all aspects of resilience. Proper interfaces between the different layers are required in order to propagate information about faults, errors, and failures in various subsystems to the subsystems that will be involved in managing them: the subsystems that need to act upon the information to contain and recover from the errors and the subsystems and will be further involved in diagnosis and repair. Furthermore, resilience techniques are often based on the assumption that a single fault will occur at a time. It is hard enough to address in a systematic manner all possible faults and practically impossible to address in a systematic manner all possible combinations of multiple faults. The ‘single fault’ assumption is statistically valid if errors are rare and are cleared rapidly. It also

requires the error-handling infrastructure to be flawless. Therefore, the correctness and the performance of the fault-handling software are paramount considerations.

6.1 Fault and error management

Each layer of running software should be able to optionally specify its dependencies, namely, which errors in other subsystems may affect it and the designated error handlers for different types of error, whether internal or external. Operational error-handling may also dump local data in support of later fault management activities (diagnosis and repair). In general, the invocation of error handlers must be carefully ordered. For simplicity, let us consider each higher-layer (or procedurally deeper) error handler as being pushed onto a stack. When passing error-handling control to successive error handlers, the system will invoke the topmost handler on the stack. When returning, the handler will indicate whether the error was successfully handled. If it was not, the next handler on the stack is invoked. In this way, error-handling passes from the most specific to the most general handler, with increasingly general actions attempted to recover from the error. The problem is complicated by the existence of a horizontal as well as a vertical organization: the error handler can be invoked on a node different from the node that signaled the error; the error can be signaled in a place different from the place where it occurred; and errors may be signaled multiple times, through different mechanisms. For example, the failure of a node can lead to an error being signaled through the hardware-monitoring infrastructure to the system console; it may cause communication timeouts, generating error messages at other nodes that communicate through the failed node; and it may generate a timeout on a system or application heartbeat. We need to ensure that recovery actions are not duplicated and are properly ordered.

As information on faults or errors propagate through the system, it is also important to properly map their semantics from level to level, into terms meaningful to each level and to the recovery abilities of each level. For example, if a bit switched in memory, the hardware layer will want to know the physical address of the affected location and will want to further localize the failure to a hardware subcomponent, such as CPU, cache, or memory. The system layer will want to know how far the error could have propagated; the application level will want to know which variables may be corrupted; and so on. Therefore, it is useful to define at each level the set of conditions that can be signaled, so that a fairly generic, portable error interface can be used to program error handlers at each level. Having such a generic classification of error types for applications will allow a more portable programming

model and a simpler evaluation of the effects of errors on application execution.

Diagnosis and repair may involve more elaborate actions that have to be coordinated across layers. For example, a node failure is recovered by replacing the node (possibly involving the global resource manager), updating routing tables and MPI structures, and restarting from the last checkpoint. Later diagnosis and recovery actions may include running detailed diagnostics and replacing the node.

A viable model for diagnosis and repair could be a software repository that allows subscriptions to fault management updates, thus allowing arbitrarily complex recovery and repair actions. In addition, these actions need information about static and dynamic configuration of the system: what the hardware and system configuration is, which applications run on which nodes, what the software configuration of the application (source code, compiler versions, library versions, etc.) was, and so on. Today, this information is typically distributed across multiple databases or is not captured at all. As a result, root cause analysis is much more painful than it should be. All configuration changes should be captured and configuration information stored in a repository, using schemata that reflect the logical system organization.

6.2 Reporting of software-detected errors

The various software layers, including the top-level application, can detect errors that were not caught by the lower-level layers. The application code may detect outliers, for example, that may indicate SDC. Therefore, reporting can also move information downward. This approach is complicated, however, since the information cannot be fully trusted (is the algorithm sure that an SDC happened, or could this be the effect of a data race?) and the information comes with less detail than information produced by lower-level detection mechanisms (the algorithm may not know where and when a bit was flipped). The passing of such information is likely to invoke a complex procedure that evaluates the reliability of the information, based on other information available to the recipient (e.g. information about the sender) and triggers activities to isolate and diagnose potentially faulty components. Such a level of activity is probably more appropriate as part of diagnosis and repair, when more complex, trainable logic can be used.

6.2.1 Error management: Algorithm hints and watchpoints.

Different parts of the software stack typically have different capabilities for handling a propagated fault or error. For example, in many situations, an application and its runtime may be able to validate its results and recover from an underlying error or fault. Similarly, a communication library may be able to establish a

dynamic alternative connection on detection of a lost communication error. In such cases, an interface that allows the different software to express their inherent fault-tolerance capabilities would be useful. Algorithmic hints would also allow lower-level software to understand what level of error semantics is useful to the upper application layers and what level of fault information could be conveyed to them. We explore algorithm hints in greater detail in Sections 6.5 and 6.6.

In many situations, the application or the lower-level subsystem can recover from an underlying error but needs to execute recovery code for that purpose. For example, an application may tolerate a corrupted data value by interpolating a replacement value from neighbor points in a mesh. This approach is most efficient if the error is detected and the recovery action enacted as soon as possible after the error occurred. In this scenario, the application (or underlying runtime) will register its ability to handle some types of error and will register the exception handler to be invoked when such an error occurs. For example, the algorithm could identify memory regions that it wants to ‘watch’ along with the recovery procedure for errors in this region. Then when an error is found by the hardware and translated up the software stack, it will trigger the appropriate exception handler, passing to it the location and type of error. The compiler and runtime need to ensure that the granularity of error reporting by hardware (e.g. ECC block) matches the granularity of software objects.

6.2.2 Error management: Communication errors. Communication errors require added attention because their effect can be global. A misrouted message could corrupt state at any node in the system. This can be handled in a variety of ways.

We can provide sufficient levels of error-handling in hardware to ensure that communication errors are fail-stop errors, where the communication fails but no incorrect message is delivered. An end-to-end protocol (a variant of the sliding window protocol) can ensure that message deletions are detected and corrected for point-to-point communication channels. However, the support for correction may require additional buffering space (to save message copies) and additional latency (to receive acknowledgments and ensure that messages will be transmitted in the right order, even in the face of errors). The problem is harder to manage for collective communications or one-sided communications. Application hints that relax message-passing semantics (e.g. relax ordering requirements) could be used to improve communication performance.

6.3 Responding to and handling of faults/errors

Various components of a system (whether software or hardware) can receive information about a fault or

error occurring in a specific part of the system. Several of these components could independently be interested in handling this fault and initiating a recovery section. These different recovery actions may be interdependent: they need to occur in a correct order, and the recovery procedure for a component may depend on the outcome of previous recovery procedures. For example, a partition may require a new node to replace a failed node. The application recovery could follow different paths if the request succeeded or failed.

Response prioritization is an inherent part of response negotiation. The system-wide resilience infrastructure will need to support mechanisms that will allow declaring response priorities of various components for the variety of faults that they might receive. In addition, interfaces are needed to allow components to specify the outcome of their recovery procedure. On failure of failure handling by the first component, the infrastructure should be able to delegate the responsibility to the next component on the list. Response negotiation can become more complicated when response priorities for components differ between various job executions.

6.4 Fault/error propagation and security implications

While a large amount of information related to faults and errors is present on systems, not all of it can be made available to all consumers because of security reasons. Often, low-level hardware/system fault information is made available only to administrators of the system. A system-wide fault-sharing infrastructure needs to have mechanisms and interfaces to control access to different types of information, for example by using capability-based security.

6.5 Top-down view of errors

Higher-level algorithms may not require notification or recovery from certain types of error, since the normal course of computation will overcome the error. A prototypical example is solving a nonlinear system of equations using Newton’s method. The basic steps of Newton’s method are to compute the residual, solve a linear system of equations with the residual on the right-hand side to obtain a direction, determine a step length along the direction, and update the iterate. For well scaled problems, Newton’s method can ignore errors in many parts of the computations without suffering ill effects.

We will often tolerate errors in the least significant digits of the mantissa of floating point numbers (say, the last eight digits), as these would be analogous to rounding errors, but we would need to detect and correct errors in the sign, exponent, or most significant mantissa digits. The recovery may be recomputation in

the case of the residual or switching to using the steepest descent direction in the step-length computation. We cannot tolerate errors in the sparsity structure of the Jacobian matrix, but we can easily tolerate low-order errors in the nonzero values of the Jacobian matrix. High-order errors in the sign, exponent, or most significant digits can also be tolerated, but there are consequences, such as degradation in the convergence rate of Newton's method and requiring extra iterations to converge to an acceptable solution. As long as the Jacobian matrix is correct or suffers from only low-order errors a large fraction of the time, then there is little impact on performance, and one may need to perform only a single extra iteration. We do note that a change in the sign or magnitude of the values can result in a positive definite matrix becoming indefinite and thus impacting linear solvers, such as the conjugate gradient method. Errors in the computation of the norm of the residual can be tolerated in the step-length calculation. Such errors may not be tolerated in the convergence tests, however, if the error results in a smaller norm of the residual that triggers premature termination of the algorithm.

Similar observations can be made about other types of algorithms and about other software subsystems. These suggest the need for an interface that enables an application or a software subsystem to provide hints and describe which lower-level errors it can tolerate. The lower software layers and the hardware could then provide differentiated levels of resilience, protecting state that the application cannot repair, if corrupted. These could include using more resilient memory, duplicating critical computation (done automatically by the compiler and runtime), or checking double-precision calculations with (cheaper) single-precision ones. Providing increasing levels of resilience would come at higher costs, with tradeoffs of both power and performance, thus requiring that the provided set of interfaces be expressive enough to allow upper-layer software to specify their tradeoff preferences.

6.6 Bottom-up view of errors

This section focuses on issues related to exposing error semantics upstream (to higher-level libraries or applications), the amount of information to be exposed, and the information to expose. Cost is a big challenge in detecting and correcting errors in the underlying hardware. The challenge is how to minimize the power and performance costs of highly effective error detection. Can we make use of high-level (e.g. application level or high-level system stack) information to minimize this cost?

Error semantics and translation. While an error can influence multiple layers of the hardware/software stack, how an error is interpreted can differ for each layer of the stack. For example, a fault on fan #1508

might be relevant at the hardware layer to correct or work around, but it might not have much semantic meaning for the application. Similarly, the fact that memory variable 'X' is corrupted might be relevant for an application, but it might not have much semantic meaning for the hardware developer, unless the virtual memory address and eventually the appropriate physical memory address translation are known.

Amount and type of error information exposed. The amount of error information propagated to upper layers needs to be tunable. While some upper layers can benefit from having information on every ECC error (corrected, or detected but not corrected) that the hardware encounters, other upper layers might be interested only in uncorrected errors. Similarly, an application might not necessarily care about errors on all of its memory regions. For example, as discussed in Section 6.5, if a higher-level library can correct memory faults on a region of memory, it might not care about the lower level of the stack returning errors for that region of memory. Such a model should also allow software architects to define the contract or expectations they have from the lower layers of the stack.

How can such hints on criticality be generated? Does the hardware need to provide low-level information to the higher layers so that the critical hints can be generated? Once hints are generated and passed down, several opportunities can exist.

Example 1. The built-in soft-error resilience (BISER) technique (Zhang et al., 2006; Mitra et al., 2007) can be configured, during system operation, to operate in one of two modes: an error-resilient mode in which BISER protection is turned on, and an economy mode in which BISER protection is turned off. Such configurability can be implemented in hardware and may be activated through software orchestration. It can minimize the system-level power cost of BISER by turning on the error-resilient mode only for critical computation. However, dynamic reliability management across multiple abstraction layers and orchestration of information flow across abstraction layers to utilize such configurability during system operation are open research questions. For BISER, one can piggyback on existing scannable signals available on-chip, but a general question concerns the costs that are incurred for such configurability at the hardware level. Can such configurability be implemented for arbitrary techniques (e.g. easy for core/thread duplication)? Is it easy for inline checking techniques such as parity prediction? What is the level of configurability that should be supported?

Example 2. One can combine software-level error resilience techniques with circuit-level techniques using a 'temporal combination' approach.

For a memory controller unit (MCU) in a multicore system on a chip (SoC), for example, we can start with

request duplication with BISER flip-flops in economy mode. We then switch the BISER flip-flops into error-resilient mode (i.e. incurring high power costs) and turn off request duplication when the systems stalls because of pending requests (which indicate high-traffic situations). We switch back to request duplication with BISER flip-flops switched to economy mode when all queues have only a few entries (to indicate low-traffic situations). Such ‘temporal combination’ simultaneously incurs a very small performance cost (performance impact similar to that of BISER-only and far better than request-duplication-only) and a small energy cost (similar to request-duplication-only and far better than BISER).

Example 3. Depending on workload, temperature sensors, and so forth, the fault-sharing framework can pass on the information to hardware to initiate fault management, for example online circuit failure prediction through reactive online self-test and diagnostics. This approach minimizes any side effects and can initiate proactive self-repair.

7 Possible scenarios

We present in this section several possible scenarios for handling failures at exascale, describe their pros and cons, and discuss technologies needed to support each scenario.

7.1 Base scenario

In the base scenario, errors are handled the same way they are handled now: applications use global

checkpoint/restart, and system software is either restarted upon failure or handles its own recovery. The obvious advantage of this scenario is that it requires (almost) no change in current application codes and requires no changes in the overall infrastructure for error recovery. (One required change will be more frequent checkpoints; with high-frequency checkpoints, it is unlikely that checkpoints will be identical to the output that goes to long-term storage or to in situ analysis.)

The performance of global checkpoint/restart schemes has been analyzed by multiple authors (Young, 1974; Daly, 2006). We recapitulate the analysis in Appendix A. This analysis enables us to compute an optimal checkpoint interval, given checkpoint time and MTTF; next we can compute the *utilization* of such systems, namely, the fraction of total computer time that is usefully applied to computation, rather than used for checkpointing and restart or wasted because of failures.

We plot in Figure 5 utilization as function of checkpoint time and recovery time. Utilization depends on the length of checkpoint and recovery relative to MTTF; if all three parameters are increased or decreased by the same ratio, then utilization is unchanged. Therefore, we express checkpoint time and recovery time as a fraction of MTTF. Figure 6 shows the same data, in the form of a contour map.

Suppose we want to achieve a utilization of more than 80%. Then Figure 6 indicates that we need to keep checkpoint time at 1%–2% of MTTF and recovery time at 2%–5% at MTTF. Assume that the MTTF of an exascale system is 30 min. Then global checkpoint




Figure 5. System utilization as a function of checkpoint and recovery time.




Figure 6. System utilization as a function of checkpoint and recovery time: contour map.

should be done in less than 20 s, and recovery in about a minute. It does not seem feasible to checkpoint so fast on disk, but it is feasible to checkpoint in a few seconds in RAM. Several schemes have been proposed for hybrid or multilevel checkpointing where frequent checkpoints are done in memory or, less frequently, on disk (Moody et al., 2010; Bautista-Gomez et al., 2011b). One can use either nonvolatile RAM to store checkpoints or volatile RAM with a RAID scheme that allows recovery from the failure of one (or more) nodes. (The first option may be constrained by the limited number of write cycles supported by various NVRAM technologies.)

Such a scheme has an obvious cost: the need to significantly increase the amount of memory by, say, 50%. This will have a significant impact on the system acquisition cost. Note, however, that the increase in power consumption is negligible. This is obvious for NVRAM but true also for DRAM, since checkpoint memory would be in standby mode most of the time.

The two main obstacles to this approach are the need to detect errors in a timely manner and the need for fast recovery. We looked in Section 3 at soft errors due to particle strikes and estimated that current technologies could be used to keep their frequency at current levels at a cost of <20% additional silicon and power. However, these numbers involve considerable uncertainty. Particle strikes are only one of multiple potential new sources of errors, and the impact of near-threshold logic was not taken into account. Furthermore, for reasons explained in Section 3.5, there is no certainty that the market will produce the low-energy, high-resilience components that would be needed to avoid silent errors

in hardware at an acceptable price. If silent errors can propagate into checkpoints, then checkpoints are not of much use.

While the time for backward recovery from checkpoint at the application level is essentially gated by I/O rates, the time for forward recovery that reboots or repairs various system software components is gated by the computation overhead of boot or repair code. Boot time of large systems may currently exceed 30 min; without a change, the boot time of an exascale supercomputer could exceed its MTBF: not a sustainable situation. This last problem is common to all envisaged scenarios for resilience. Therefore, advances that reduce boot time and repair time for the system infrastructure at exascale are essential. Also essential are advances that reduce the likelihood of system failures, in particular software failures.

7.2 System software scenario

In the second scenario, hardware is assumed not to provide enough detection, and therefore SDC events occur too frequently to be ignored. Instead, we assume that data corruption can be prevented, detected, and corrected or else tolerated with no change to the application software.

Not all hardware errors have the same severity. A bit flip in a large array of data may have little impact on the final answer; a bit flip in a program counter or a data pointer is likely to have a stranger, less predictable impact; and a bit flip in a page table or a routing table is likely to have a catastrophic impact. Luckily, the software error detection schemes described in Section 5.4.1 are more likely to detect the ‘bad errors’: those that will have a significant impact on the final answer or will cause a crash. Furthermore, redundancy can be used in order to reduce the probability of ‘bad errors’. Critical computations can be executed twice (and the redundancy can be introduced automatically by a compiler; see Reis et al., 2005b; Yu et al., 2009); more reliable memory may be used for more sensitive data, and so forth.

A plausible hypothesis is that silent hardware errors fall into two categories: ‘pleasant errors’ that can be treated as aleatoric uncertainty in the computation and ‘nasty errors’ that, essentially, change the computation model. The latter must be treated as epistemic errors that cannot be modeled as statistical noise and have to be avoided or corrected. Fortunately, ‘nasty errors’ are likely to be less frequent than ‘pleasant errors’ in large scientific codes and are easier to avoid or correct. If this hypothesis is correct, then SDC events can be survived with little to no change in application codes. This hypothesis needs to be validated for all or a large fraction of large scientific workloads.

The system scenario also covers schemes for using local restart, thus reducing restart overhead, provided

that the construction of node clusters (application containers) can be automated.

7.3 Application scenarios

Handling resilience without changes in application codes may turn out to be too expensive. We envisage two subcases: those in which the application code has to handle only tolerance or detection and those in which the application code also has to handle correction.

We discussed fault-tolerant algorithms in Section 5.3, algorithmic fault detection in Section 5.4.2, and algorithmic recovery in Section 5.6.4. The main issue with these techniques is that they are specific to one or to a family of algorithms. We need generic techniques that will apply to *all* computations of interest for the exascale era and efficient techniques that will apply to the large majority of these computations.

Another issue is how to compose different approaches to resilience. If one module can tolerate silent bit flips, another module can detect them efficiently and recover using checkpoints, and yet another module needs redundant execution, how are these three modules coupled in one application?

8 Suggested actions

We outline in this section actions that are suggested by this workshop.

8.1 Information gathering

The different scenarios imply very different strategies for achieving the required level of resilience, from possibly significant investments in hardware that has little use outside extreme-scale computing to possibly significant investments in recoding existing applications. At this time, we do not have enough information to choose a direction; more information-gathering is essential. We propose several activities for that purpose.

8.1.1 Characterization of sources of failures on current systems. DOE has a rich source of information in the form of the message logs that are collected at each of the supercomputing centers at DOE labs. Unfortunately, most of this data is not centrally collected; also, different vendors use distinct terminologies, so that data cannot be directly compared. To the best of our knowledge, there are no vendor restrictions on the publication of data owned by the various centers. Initial discussions with vendors indicate a willingness to help analyze the data.

We propose to establish, as soon as possible, a centralized repository within DOE that will systematically collect event logs and other relevant information from

all DOE supercomputing centers. In parallel, we propose to invest in tools to normalize these logs into a vendor-neutral notation and to anonymize them. DOE would then make these cleansed logs available to the broader research community.

We note that the paper of Schroeder and Gibson on ‘Understanding failures in petascale computers’ (Schroeder and Gibson, 2010) cites three repositories for computer failure data. Two (at Los Alamos National Laboratory and NERSC) do not seem to be accessible on the web. The third, the Computer Failure Data Repository (CFDR) at <http://cfdr.usenix.org>, which is maintained by Bianca Schroeder, is easily accessible. This situation suggests that a community effort will be more productive than the individual efforts of supercomputing centers.

Event logs provide failure symptoms but do not provide a root cause for each failure. Root cause analysis is now a tedious manual process that engages much of the time of the staff at supercomputing centers. We propose two efforts on root cause analysis:

1. Develop a registration system that will facilitate recording the results of the manual root cause analysis. The goal is to annotate event logs with the results of such analyses.
2. Develop better tools for root cause analysis. Existing software products, such as SMARTS of EMC (EMC, 2014), could be a good start for such development.

8.1.2 Study of frequency of silent errors. Currently there exists a large uncertainty about the frequency of SDC events. On the one hand, the practice of supercomputer users is to assume such events do not occur. On the other hand, anecdotal evidence on the nonreproducibility of computations that are supposed to be bit-reproducible suggests they do occur, and occur quite frequently.

We propose to push a study on the frequency of SDC events on current supercomputers. Such a study could be effected by running a background job on as many nodes as possible on various supercomputers. The job would produce bit-reproducible, testable results and be used to detect SDCs.

8.1.3 Refinement of estimates on future hardware technologies. The main uncertainty about future roadblocks to resilience concerns the frequency of hardware SDC events. Our analysis showed that cosmic-radiation-induced SDCs could be managed at a cost of less than 20% in circuitry and in power consumption, using current methods. More research in this area could further reduce the gap. However, the study ignored other issues (subthreshold logic, aging). In any case, the main uncertainty about future hardware technologies is less

about what can be done and more about what will be done by industry, given market forces. It will be useful to complement technological studies with economic studies, based on the evolution of different markets (high-end server, cloud, mobile). The key question to be addressed is the following: what is the market size for processors that have low power, high resilience, and high floating-point performance?

8.2 Research areas

We divide research directions into three categories.

Necessary technologies: Technologies that will be necessary for resilience at extreme scale, no matter what scenario ends up being pursued.

Generally useful technologies: Technologies that will be useful no matter what scenario ends up taking effect.

Scenario-specific technologies: Technologies that will come into play only under a subset of the scenarios.

DOE investments in R&D should focus on the roadblocks we know will certainly exist, and less so on roadblocks that are still hypothetical. On the other hand, one may justify investments in scenario-specific technologies as a risk-reduction action, if the technology is necessary under some plausible scenario and the time lag from research to deployment is expected to be significant.

8.2.1 Necessary technologies. In any scenario, it will be essential to reduce the frequency of system failures, contain them, and reduce recovery time from system failures. Some of the problems may have simple engineering solutions, for example, fast boot from NVM. Solutions to other problems may require new structures and mechanisms for global system services. Some of the current research on error containment that is now focused on application errors could be fruitfully applied to system errors. Faster recovery from file system failures will be important.

Another critical technology is the communication infrastructure that enables recovery actions at different levels of the system. This infrastructure will need to be as resilient as the current out-of-band networks that collect hardware-monitoring information and channel it to the hardware-monitoring console. But the infrastructure also will need to handle software failures and avoid the sequential bottleneck of one global monitoring point.

8.2.2 Generally useful technologies. Some technologies are useful no matter what scenario takes effect. One example is fault prediction and avoidance: predicting node failures and migrating a node workload before the node fails. Successful fault prediction and avoidance

effectively increase the system MTBF, thus increasing the system utilization.

Another example is provided by technologies for fault containment. Avoiding a global restart can reduce the time and energy consumed by restarts, thus improving system performance.

8.2.3 Scenario-specific technologies. Scenario-specific technologies include all the technologies that will be required if SDCs become a major problem: technologies for system software error detection, containment, and correction, and technologies for application-level error tolerance, detection, containment, and correction.

Arguably, the choice between handling errors in hardware or in firmware is a vendor choice. Vendors will choose one or the other, or a mix of the two, according to the relative non-recurring and recurring costs of the two approaches. Research in DOE can help in exploring firmware-level resilience solutions. We recommend a co-design collaboration between DOE research and vendors in exploring the right mix of hardware and system software approaches that would provide the appearance of a failure-free system to the application layer.

Application-level error-handling is a much more significant departure from current practice, one that should be entertained only if the other options are not feasible or have a significant cost. Application-level error correction will require new services from the underlying hardware and software, for example the ability to provide differentiated resilience quality for computations or storage, fault-tolerance at the level of MPI and other global libraries, and mechanisms for signaling errors to application code. Since these are needed for research in application-level error-handling, their development should be a priority.

A main focus on application-level error-handling should be on generic techniques that apply to all applications or large classes of applications. These are needed in order to avoid having to develop a unique solution for each application code. If different techniques are used for different codes, then one will need methods for composing these techniques.

We note that application-level tolerance or detection of SDCs is more important than application-level correction, since global/checkpoint restart is still viable at exascale, provided one can ignore or detect errors.

8.3 Integration

Much of the current research on resilience is addressing small sections of the problem, for example how to tolerate or detect SDC errors for a particular algorithm. Point solutions are useful only if they fit in an overall resilience architecture. For example, algorithm error-handling may assume that some system services

continue to be available after an error has occurred and may be able to handle some errors (a bit flip in data) while ignoring other errors (a bit flip in a pointer). These assumptions and limitations must be made explicit in order to ensure that error modes ignored by the point solution are either sufficiently rare or handled by another point solution.

We need to develop a *resilience architecture* that specifies (1) which errors are assumed to occur and which errors are assumed to be so rare as to be ignored and (2) what the division of labor is between the various layers of the system in handling frequent errors.

As long as we have not converged to one scenario, we will have multiple resilience architectures. But each of them must be brought to a reasonable level of completeness in order to make sure the different approaches are comprehensive.

Acknowledgements

We thank the U.S. Department of Energy for its financial support of ICiS; the ICiS director and steering committee for the support provided to our workshop; and, in particular, Cheryl Zidel for her outstanding administrative support before, during, and after the workshop. We also thank Gail Pieper for her thorough editing of this report, and Lucy Novell and Daniel Katz for useful comments.

Funding

This work was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research (contract DE-AC02-06CH11357).

References

- Agostinelli M, Pae S, Yang W, et al. (2005) Random charge effects for PMOS NBTI in ultra-small gate area devices. In: *Proceedings of the 2005 IEEE international reliability physics symposium (IRPS)*, pp. 529–532.
- Ahn DH, Supinski BRD, Laguna I, et al. (2009) Scalable temporal order analysis for large scale debugging. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Austin TM (1999) DIVA: A reliable substrate for deep sub-micron microarchitecture design. In: *Proceedings of the annual international symposium on microarchitecture (MICRO)*, pp. 196–207.
- Avizienis A (1973) Arithmetic algorithms for error-coded operands. *IEEE Transactions on Computers* C-22(6): 567–572.
- Avizienis A, Laprie J, Randell B, et al. (2004) Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing* 1(1): 11–33.
- Avritzer A, Bondi A, Grottke M, et al. (2006) Performance assurance via software rejuvenation: Monitoring, statistics and algorithms. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*, pp. 435–444.
- Bailey FR, Bell G, Blondin J, et al. (2007) Petascale metrics panel report. Available at: http://research.microsoft.com/en-us/um/people/gbell/supers/ascac_petascale_metrics_panel_report_and_executive_summary_2007-02-12.pdf (accessed 25 February 2014)
- Ballesteros FJ, Evans N, Forsyth C, et al. (2012) Nix: A case for a manycore system for cloud computing. *Bell Labs Technical Journal* 17(2): 41–54.
- Banerjee P and Abraham J (1986) Bounds on algorithm-based fault tolerance in multiple processor systems. *IEEE Transactions on Computers* C-35(4): 296–306.
- Banerjee P, Rahmeh J, Stunkel C, et al. (1990) Algorithm-based fault tolerance on a hypercube multiprocessor. *IEEE Transactions on Computers* 39(9): 1132–1145.
- Bautista-Gomez LA, Tsuboi S, Komatitsch D, et al. (2011a) FTI: High performance fault tolerance interface for hybrid systems. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Bautista-Gomez L, Komatitsch D, Maruyama N, et al. (2011b) FTI: High performance fault tolerance interface for hybrid systems. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Birge J and Louveaux F (1997) *Introduction to Stochastic Programming*. Berlin: Springer Verlag.
- Bland W, Bouteiller A, Herault T, et al. (2012) An evaluation of user-level failure mitigation support in MPI. In: Träff J, Benkner S and Dongarra J (eds) *Recent Advances in the Message Passing Interface*. New York, NY: Springer, pp. 193–203.
- Borkar S (2005) Designing reliable systems from unreliable components: The challenges of transistor variability and degradation. *IEEE Micro* 25(6): 10–16.
- Bosilca G, Delmas R, Dongarra J, et al. (2009) Algorithm-based fault tolerance applied to high performance computing. *Journal of Parallel and Distributed Computing* 69(4): 410–416.
- Bouteiller A, Herault T, Bosilca G, et al. (2011) Correlated set coordination in fault tolerant message logging protocols. In: *Euro-Par 2011: Parallel Processing Workshops* (eds E Jeannot, R Namyst and R Jean), 29 August– 2 September 2011, France, pp. 51–64. New York, NY: Springer.
- Bouteiller A, Herault T, Krawezik G, et al. (2006) MPICH-V project: A multiprotocol automatic fault-tolerant MPI. *International Journal of High Performance Computing Applications* 20(3): 319–333.
- Bower F, Sorin D and Ozev S (2007) Online diagnosis of hard faults in microprocessors. *ACM Transactions on Architecture and Code Optimization* 4(2).
- Bronevetsky G, Laguna I, Bagchi S, et al. (2010) AutoMaDeD: Automata-based debugging for dissimilar parallel tasks. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*, pp. 231–240.
- Cai K and Qin Z, Memory Device with Soft-Decision Decoding. US Patent 20130107611 A1, May 2, 2013.
- Cappello F, Geist A, Gropp B, et al. (2009) Toward exascale resilience. *International Journal of High Performance Computing Applications* 23(4): 374–388.
- Cappello F, Guermouche A and Snir M (2010) On communication determinism in parallel HPC applications.

- In: *Proceedings of the 19th international conference on computer communications and networks (ICCCN)*, pp. 1–8.
- Carulli J and Anderson T (2005) Test connections-tying application to process. In: *IEEE Workshop on Silicon Errors in Logic–System Effects*, Stanford University, CA.
- Castelli V, Harper RE, Heidelberger P, et al. (2001) Proactive management of software aging. *IBM Journal of Research and Development* 45(2): 311–332.
- Chan JTY, Tseng CW, Chu YC, et al. (1998) Experimental results for IDDQ and VLV testing. In: *Proceedings of the IEEE VLSI test symposium*, pp. 118–125.
- Chen D, Eisley NA, Heidelberger P, et al. (2011) The IBM Blue Gene/Q interconnection network and message unit. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Chen Z and Dongarra J (2006) Algorithm-based checkpoint-free fault tolerance for parallel matrix computations on volatile resources. In: *Proceedings of the 20th international parallel and distributed processing symposium (IPDPS)*.
- Chow P (2007) *Stochastic Partial Differential Equations*. Boca Raton/London/New York: Chapman & Hall/CRC.
- Chung J, Lee I, Sullivan M, et al. (2012) Containment domains: A scalable, efficient, and flexible resilience scheme for exascale systems. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Conn AR, Gould NI and Toint PL (1987) *Trust-Region Methods*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Daly J, Adolf B, Borkar S, et al. (2012) Inter agency workshop on HPC resilience at extreme scale. Available at: <http://institutes.lanl.gov/resilience/docs/Inter-AgencyResilienceReport.pdf> (accessed 25 February 2014).
- Daly JT (2006) A higher order estimate of the optimum checkpoint interval for restart dumps. *Future Generation Computer Systems* 22(3): 303–312.
- Dean J and Ghemawat S (2008) MapReduce: Simplified data processing on large clusters. *Communications of the ACM* 51(1): 107–113.
- DeBardeleben N, Laros J, Daly J, et al. (2010b) High-end computing resilience: Analysis of issues facing the HEC community and path-forward for research and development. Technical Report LA-UR-10-00030, DARPA, VA. available at <http://www.csm.ornl.gov/~engelmann/publications/debardeleben09high-end> 2/25/14
- DeHon A, Carter N and Quinn H (eds) (2011) Final report for CCC cross-layer reliability visioning study. 3 March Available at: <http://xlayer.org/FinalReport> (accessed 25 February 2014).
- Dimitrov M and Zhou H (2007) Unified architectural support for soft-error protection or software bug detection. In: *Proceedings of the conference on parallel architecture and compilation techniques*, pp. 73–82.
- Dixit A, Heald R and Wood A (2009) Trends from ten years of soft error experimentation. In: *The workshop on silicon* Available at: http://softerrors.info/selse/images/selse_2009/Papers/selse5_submission_29.pdf (accessed 25 February 2014).
- Dongarra J, Beckman P, Moore T, et al. The international exascale software project roadmap International Journal of High Performance Computing Applications, 25(1), 3–60, 2011.
- Downing R, Nowak J and Tuomenoksa L (1964) No. 1 ESS maintenance plan. *Bell System Technical Journal* 43(5): 1961–2019.
- Du P, Bouteiller A, Bosilca G, et al. (2012) Algorithm-based fault tolerance for dense matrix factorizations. In: *Proceedings of the 17th ACM SIGPLAN symposium on principles and practice of parallel programming*, New York, NY, pp. 225–234.
- Elnozahy ENM, Alvisi L, Wang YM, et al. (2002) A survey of rollback-recovery protocols in message-passing systems. *ACM Computing Surveys* 34(3): 375–408.
- Elnozahy (editor) System Resilience at Extreme Scale White Paper available at <http://citeseerx.ist.psu.edu/viewdoc/download?rep=rep1&type=pdf&doi=10.1.1.205.4240accessed2/25/14>
- EMC (2014) Smarts: Automated IT management enabling service assurance. Available at: <http://www.emc.com/it-management/smarts/index.htm> (accessed 25 February 2014).
- Ernst MD, Perkins JH, Guo PJ, et al. (2007) The Daikon system for dynamic detection of likely invariants. *Science of Computer Programming* 69(1): 35–45.
- Fadden S (2012) An introduction to GPFS version 3.5. Available at: www-03.ibm.com/systems/jo/resources/introduction-to-gpfs-3-5.pdf (accessed 25 February 2014).
- Fagg G and Dongarra J (2000) FT-MPI: Fault tolerant MPI, supporting dynamic applications in a dynamic world. In: Dongarra J, et al. (eds) *Recent Advances in Parallel Virtual Machine and Message Passing Interface* (Lecture Notes in Computer Science, vol. 1908). Berlin/Heidelberg: Springer, pp. 346–353.
- Feng S, Gupta S, Ansari A, et al. (2010) Shoestring: Probabilistic soft error reliability on the cheap. In: *Proceedings of the international conference on architectural support for programming languages and operating systems (ASPLOS)*, pp. 385–396.
- Ferreira KB, Stearley J, Laros JH III, et al. (2011) Evaluating the viability of process replication reliability for exascale systems. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Fletcher R (1981) *Practical Methods of Optimization*. Volume 2: Constrained Optimization. New York, NY: John Wiley & Sons.
- Fujita H, Schreiber R and Chien AA (2013) It's time for new programming models for unreliable hardware. In: *Proceedings of the international conference on architectural support for programming languages and operating systems (ASPLOS)*.
- Gainaru A, Cappello F and Kramer W (2012a) Taming of the shrew: Modeling the normal and faulty behavior of large-scale HPC systems. In: *Proceedings of the IEEE international parallel & distributed processing symposium (IPDPS)*.
- Gainaru A, Cappello F, Fullop J, et al. (2011a) Adaptive event prediction strategy with dynamic time window for large-scale HPC systems. In: *Proceedings of managing large-scale systems via the analysis of system logs and the application of machine learning techniques (SLAM'11)*, pp. 4:1–4:8.
- Gainaru A, Cappello F, Snir M, et al. (2012b) Fault prediction under the microscope: A closer look into HPC

- systems. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Gainaru A, Cappello F, Trausan-Matu S, et al. (2011b) Event log mining tool for large scale HPC systems. In: *Euro-Par 2011: Parallel Processing Workshops*. New York M Alexander, P D'Ambra, A Belloum, et al. (eds), NY: Springer.
- Gao B, Zhang H, Chen B, et al. (2011) Modeling of retention failure behavior in bipolar oxide-based resistive switching memory. *IEEE Electron Device Letters* 32(3): 276–278.
- Gao Q, Qin F and Panda DK (2007) DMTracker: Finding bugs in large-scale parallel programs by detecting anomaly in data movements. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Gao Q, Zhang W and Qin F (2010) FlowChecker: Detecting bugs in MPI libraries via message flow checking. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Gattiker A, Nigh P, Grosch D, et al. (1996) Current signatures for production testing [CMOS ICs]. In: *IEEE international workshop on IDDQ testing*, pp. 25–28.
- Geist A, Lucas B, Snir M, et al. (2012) U.S. Department of Energy fault management workshop. Technical report, U.S. Department of Energy, DC.
- Gill B, Seifert N and Zia V (2009) Comparison of alpha-particle and neutron-induced combinational and sequential logic error rates at the 32nm technology node. In: *IEEE international reliability physics symposium*, pp. 199–205.
- Golubeva O, Rebaudengo M, Reorda MS, et al. (2003) Soft-error detection using control flow assertions. In: *Proceedings of the international symposium on defect and fault tolerance in VLSI systems*, pp. 581–588.
- Griewank A and Corliss G (1991) *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Grottke M and Trivedi KS (2007) Fighting bugs: Remove, retry, replicate, and rejuvenate. *IEEE Computer* 40(2): 107–109.
- Guermouche A, Ropars T, Brunet E, et al. (2011) Uncoordinated checkpointing without domino effect for send-deterministic MPI applications. In: *IEEE international parallel & distributed processing symposium (IPDPS)*, pp. 989–1000.
- Guermouche A, Ropars T, Snir M, et al. (2012) HydEE: Failure containment without event logging for large scale send-deterministic MPI applications. In: *IEEE international parallel & distributed processing symposium (IPDPS)*, pp. 1216–1227.
- Gunnels J, Katz D, Quintana-Orti E, et al. (2001) Fault-tolerant high-performance matrix multiplication: Theory and practice. In: *Proceedings of the international conference on dependable systems and networks (DSN)*, pp. 47–56.
- Hackbusch W (1985) *Multi-Grid Methods and Applications*. Berlin: Springer-Verlag.
- Hafner JL, Deenadhayalan V, Belluomini W, et al. (2008) Undetected disk errors in RAID arrays. *IBM Journal of Research and Development* 52(4.5): 413–425.
- Hamming R (1987) *Numerical Methods for Scientists and Engineers*. New York: Dover Publications.
- Hangal S and Lam MS (2002) Tracking down software bugs using automatic anomaly detection. In: *Proceedings of the 2002 international conference on software engineering*.
- Hao H and McCluskey E (1993) Very-low-voltage testing for weak CMOS logic ICs. In: *Proceedings of the IEEE international test conference (ITC)*, pp. 275–284.
- Hari SKS, Adve SV and Naeimi H (2012a) Low-cost program-level detectors for reducing silent data corruptions. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*.
- Hari SKS, Adve SV, Naeimi H, et al. (2012b) Relyzer: Exploiting application-level fault equivalence to analyze application resiliency to transient faults. In: *Proceedings of the international conference on architectural support for programming languages and operating systems (ASPLOS)*.
- Hari SKS, Li ML, Ramachandran P, et al. (2009) mSWAT: Low-cost hardware fault detection and diagnosis for multicore systems. In: *Proceedings of the annual international symposium on microarchitecture (MICRO)*, pp. 122–132.
- Hazucha P, Karnik T, Bloechel SWB, et al. (2003) Measurements and analysis of SER tolerant latch in a 90 nm dual-Vt CMOS process. In: *IEEE custom integrated circuits conference*, pp. 617–620.
- Hedges R, Loewe B, McLarty T, et al. (2005) Parallel file system testing for the lunatic fringe: The care and feeding of restless I/O power users. In: *Proceedings of the 22nd IEEE/13th NASA Goddard conference on mass storage systems and technologies*, pp. 3–17.
- Heien E, Kondo D, Gainaru A, et al. (2011) Modeling and tolerating heterogeneous failures in large parallel systems. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Heiser G, Ryzhyk L, Von Tessin M, et al. (2011) What if you could actually trust your kernel. In: *13th workshop on hot topics in operating systems (HotOS)*.
- Hess WN, Patterson HW, Wallace R, et al. (1959) Cosmic-ray neutron energy spectrum. *Physical Review* 116(2): 445.
- Hogan S, Hammond J and Chien AA (2012) An evaluation of difference and threshold techniques for efficient checkpointing. In: *2nd workshop on fault-tolerance for HPC at extreme scale (FTXS 2012)*.
- Huang KH and Abraham J (1984) Algorithm-based fault tolerance for matrix operations. *IEEE Transactions on Computers* C-33(6): 518–528.
- Hunter R (1975) Engine failure prediction techniques. *Aircraft Engineering and Aerospace Technology* 47(3): 4–14.
- Hwang AA, Stefanovici IA and Schroeder B (2012) Cosmic rays don't strike twice: Understanding the nature of DRAM errors and the implications for system design. In: *Proceedings of the international conference on architectural support for programming languages and operating systems (ASPLOS)*, pp. 111–122.
- Ibe E, Taniguchi H, Yahagi Y, et al. (2010) Impact of scaling on neutron-induced soft error in SRAMs from a 250 nm to a 22 nm design rule. *IEEE Transactions on Electron Devices* 57(7): 1527–1538.
- Katz D and Some R (2003) NASA advances robotic space exploration. *Computer* 36(1): 52–61.

- Katz DS, Daly J, DeBardeleben N, et al. (2009) 2009 fault tolerance for extreme-scale computing workshop. Technical report ANL/MCS-TM-312, Argonne National Laboratory, IL.
- Kerbyson D, Rajamony R and Van Hensbergen E (2012) Performance health monitoring for large-scale systems. In: *Second international workshop on high-performance infrastructure for scalable tools*.
- Kubota K and Iri M (1992) Estimates of rounding errors with fast automatic differentiation and interval analysis. *Journal of Information Processing* 14(3): 508–515.
- Kundu S, Mak T and Galivanche R (2004) Trends in manufacturing test methods and their implications. In: *Proceedings of the international test conference (ITC)*, pp. 679–687.
- Laguna I, Ahn DH, de Supinski BR, et al. (2012) Probabilistic diagnosis of performance faults in large-scale parallel applications. In: *Proceedings of the 21st international conference on parallel architectures and compilation techniques*, pp. 213–222.
- Laguna I, Gambin T, de Supinski BR, et al. (2011) Large scale debugging of parallel tasks with AutomaDeD. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Lange J, Pedretti K, Hudson T, et al. (2010) Palacios and kiten: New high performance operating systems for scalable virtualized and native supercomputing. In: *IEEE international symposium on parallel & distributed processing (IPDPS)*, pp. 1–12.
- Lee GL, Ahn DH, Arnold DC, et al. (2007) Benchmarking the stack trace analysis tool for Blue Gene/L. In: *International conference on parallel computing: Architectures, algorithms and applications (ParCo)*.
- Lee GL, Ahn DH, Arnold DC, et al. (2008) Lessons learned at 208K: Towards debugging millions of cores. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Li ML, Ramachandran P, Sahoo S, et al. (2008a) Trace-based microarchitecture-level diagnosis of permanent hardware faults. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*.
- Li ML, Ramachandran P, Sahoo S, et al. (2008b) Understanding the propagation of hard errors to software and implications for resilient systems design. In: *Proceedings of the international conference on architectural support for programming languages and operating systems (ASPLOS)*, pp. 265–276.
- Lindekugel K, DiGirolamo A and Stanzone D (2008) Architecture for an offline parallel debugger. In: *International symposium on parallel and distributed processing with applications (ISPA'08)*, pp. 227–235.
- Linderoth J and Wright S (2003) Decomposition algorithms for stochastic programming on a computational grid. *Computational Optimization and Applications* 24(2): 207–250.
- Lo JC (1994) Reliable floating-point arithmetic algorithms for error-coded operands. *IEEE Transactions on Computers* 43(4): 400–412.
- Lo J, Thanawastien S and Rao T (1989) Concurrent error detection in arithmetic and logical operations using Berger codes. In: *Proceedings of 9th symposium on computer arithmetic*, pp. 233–240.
- Los Alamos National Lab (2006) Operational data to support and enable computer science research. Available at: <http://institutes.lanl.gov/data/fdata/> (accessed 25 February 2014).
- Lourenço J and Cunha J (2001) Fiddle: A flexible distributed debugger architecture. In: *International conference on computational science (ICCS)*, pp. 821–830.
- Lu G, Zheng Z and Chien AA (2013) When are multiple checkpoints needed? In: *3rd workshop on fault-tolerance for HPC at extreme scale (FTXS 2013)*.
- Lunardini D, Narasimham B, Ramachandran V, et al. (2004) A performance comparison between hardened-by-design and conventional-design standard cells. In: *2004 workshop on radiation effects on components and systems, radiation hardening techniques and new developments*.
- Lyle G, Cheny S, Pattabiraman K, et al. (2009) An end-to-end approach for the automatic derivation of application-aware error detectors. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*, pp. 584–589.
- Maxwell P, O'Neill P, Aitken R, et al. (2000) Current ratios: A self-scaling technique for production IDDQ testing. In: *Proceedings of the international test conference (ITC)*, pp. 1148–1156.
- Meixner A, Bauer ME and Sorin DJ (2007) Argus: Low-cost, comprehensive error detection in simple cores. In: *Proceedings of the annual international symposium on microarchitecture (MICRO)*, pp. 210–222.
- Mirgorodskiy AV, Maruyama N and Miller BP (2006) Problem diagnosis in large-scale computing environments. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Mitchell R (1977) *The Underground Grammarian*, Vol., No. 1, January. Available at <http://www.sourcetext.com/grammarian/> (accessed 25 February 2014).
- Mitra S, Zhang M, Seifert N, et al. (2007) Built-in soft error resilience for robust system design. In: *IEEE international conference on integrated circuit design and technology*.
- Mokhtarani A, Kramer W and Hick J (2008) Reliability results of NERSC systems. <https://publications.lbl.gov/islandora/object/ir%3A150330> (accessed 25 February 2014).
- Moody A, Bronevetsky G, Mohror K, et al. (2010) Design, modeling, and evaluation of a scalable multi-level checkpointing system. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Moré JJ and Wild SM (2012) Estimating derivatives of noisy simulations. *ACM Transactions of Mathematical Software* 38(3): 19: 1–19: 21.
- MPIPlugin (2013) MPI plugin for KDevelop. Available at: <http://sourceforge.net/projects/mpiplugin/> (accessed 25 February 2014).
- Nakano J, Montesinos P, Gharachorloo K, et al. (2006) ReVive I/O: Efficient handling of I/O in highly-available rollback-recovery servers. In: *Proceedings of the international symposium on high performance computer architecture (HPCA)*.
- Naksinehaboon N, Taerat N, Leangsuksun C, et al. (2010) Benefits of software rejuvenation on HPC systems. In: *International symposium on parallel and distributed processing with applications (ISPA)*, pp. 499–506.

- Nassif S, Kleeberger V and Schlichtmann U (2012) Goldilocks failures: Not too soft, not too hard. In: *2012 IEEE international reliability physics symposium (IRPS)*, pp. 2F–1.
- NCAR (2014) Community earth system model. Available at: <http://www2.cesm.ucar.edu/> (accessed 25 February 2014).
- Network Working Group (2009) The syslog protocol. Available at: <http://tools.ietf.org/html/rfc5424> (accessed 25 February 2014).
- Nigh P and Gattiker A (2000) Test method evaluation experiments and data. In: *Proceedings of the international test conference (ITC)*, pp. 454–463.
- Oh J, Washington SP and Nam D (2006) Accident prediction model for railway-highway interfaces. *Accident Analysis and Prevention* 38(2): 346–356.
- Oliner A and Stearley J (2007) What supercomputers say: A study of five system logs. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*, pp. 575–584.
- Park Y, Van Hensbergen E, Hillenbrand M, et al. (2012) FusedOS: Fusing LWK performance with FWK functionality in a heterogeneous environment. In: *24th international symposium on computer architecture and high performance computing (SBAC-PAD)*, pp. 211–218.
- Pattabiraman K, Nakka N, Kalbarczyk Z, et al. (2008) SymPLFIED: Symbolic program-level fault injection and error detection framework. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*.
- Pattabiraman K, Saggese GP, Chen D, et al. (2006) Dynamic derivation of application-specific error detectors and their implementation in hardware. In: *European dependable computing conference*, pp. 97–108.
- Prvulovic M, Zhang Z and Torrellas J (2002) ReVive: Cost-effective architectural support for rollback recovery in shared-memo multiprocessors. In: *Proceedings of the annual international symposium on computer architecture (ISCA)*.
- Racunas P, Constantinides K, Manne S, et al. (2007) Perturbation-based fault screening. In: *Proceedings of the international symposium on high performance computer architecture (HPCA)*, pp. 169–180.
- Ramachandran P (2011) Detecting and recovering from in-core hardware faults through software anomaly treatment. PhD Thesis, University of Illinois at Urbana Champaign, IL.
- Randall A V (2006) The Eckert tapes: Computer pioneer says ENIAC team couldn't afford to fail – and didn't. *Computerworld* 40(8): 18
- Rao TRN (1974) *Error Coding for Arithmetic Processors*. Orlando, FL: Academic Press, Inc.
- Reddy V, Krishnan A, Marshall A, et al. (2005) Impact of negative bias temperature instability on digital circuit reliability. *Microelectronics Reliability* 45(1): 31–38.
- Reis G, Chang J, Vachharajani N, et al. (2005a) Software-controlled fault tolerance. *ACM Transactions on Architecture and Code Optimization* 2(4): 366–396.
- Reis GA, Chang J, Vachharajani N, et al. (2005b) SWIFT: Software implemented fault tolerance. In: *Proceedings of the international symposium on code generation and optimization*, pp. 243–254.
- Rogue Wave Software (2013) TotalView Debugger. Available at: <http://www.roguewave.com/products/totalview.aspx> (accessed 25 February 2014).
- Ropars T, Guermouche A, Uçar B, et al. (2011) On the use of cluster-based partial message logging to improve fault tolerance for MPI HPC applications. *Euro-Par 2011: Parallel Processing Workshops*. In: *17th International Euro-Par Conference* (eds J Emmanuel, N Raymond and R Jean), Bordeaux, France, 29 August– 2 September 2011, pp. 567–578. New York, NY: Springer.
- Roth PC, Arnold DC and Miller BP (2003) MRNet: A software-based multicast/reduction network for scalable tools. In: *International conference for high-performance computing, networking, storage and analysis (SC)*.
- Roy-Chowdhury A, Bellas N and Banerjee P (1996) Algorithm-based error-detection schemes for iterative solution of partial differential equations. *IEEE Transactions on Computers* 45(4): 394–407.
- Sahoo S, Li ML, Ramchandran P, et al. (2008) Using likely program invariants to detect hardware errors. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*, pp. 70–79.
- Salfner F, Lenk M and Malek M (2010) A survey of online failure prediction methods. *ACM Computing Surveys* 42: 1–42.
- Saxena N and McCluskey E (2002) Dependable adaptive computing systems – the ROAR project. In: *IEEE international conference on systems, man, and cybernetics*, pp. 2172–2177.
- Schroeder B and Gibson GA (2007) Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you. In: *Proceedings of the 5th USENIX conference on file and storage technologies (FAST)*, pp. 1–16.
- Schroeder B and Gibson GA (2010) A large-scale study of failures in high-performance computing systems. *IEEE Transactions on Dependable and Secure Computing* 7(4): 337–350.
- Schroeder B, Pinheiro E and Weber WD (2009) DRAM errors in the wild: A large-scale field study. In: *Proceedings of the eleventh international joint conference on measurement and modeling of computer systems*, pp. 193–204.
- Seltborg P, Polanski A, Petrochenkov S, et al. (2005) Radiation shielding of high-energy neutrons in SAD. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 550(1): 313–328.
- Shipman G, Dillow D, Oral S, et al. (2010) Lessons learned in deploying the world's largest scale Lustre file system. In: *The 52nd Cray user group conference*.
- Slayman C (2011) Soft error trends and mitigation techniques in memory devices. In: *Proceedings of the annual reliability and maintainability symposium (RAMS)*, pp. 1–5.
- Slegel TJ, Averill RM III, Check MA, et al. (1999) IBM's S/390 G5 microprocessor design. *IEEE Micro* 19(2): 12–23.
- Snir M and Bader DA (2004) A framework for measuring supercomputer productivity. *International Journal for High Performance Computing Applications* 18(4): 417–432.
- Sorin D, Martin MMK, Hill MD, et al. (2002) SafetyNet: Improving the availability of shared memory multiprocessors with global checkpoint/recovery. In: *Proceedings of*

- the annual international symposium on computer architecture (ISCA).*
- Spainhower L and Gregg T (1999) IBM S/390 parallel enterprise server G5 fault tolerance: A historical perspective. *IBM Journal of Research and Development* 43(5.6): 863–873.
- Sridharan V and Liberty D (2012) A study of DRAM failures in the field. In: International conference for high-performance computing, networking, storage and analysis (SC).
- Stearley J (2005) Defining and measuring supercomputer reliability, availability, and serviceability (RAS). In: *Proceedings of the Linux clusters institute conference.*
- Taleb N (2010) *The Black Swan: The Impact of the Highly Improbable.* New York: Random House Trade Paperbacks.
- Trottenberg U, Oosterlee C and Schüller A (2001) *Multigrid.* New York, NY: Academic Press.
- Turmon M, Granat R and Katz D (2000) Software-implemented fault detection for high-performance space applications. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*, pp. 107–116.
- Turmon M, Granat R, Katz D, et al. (2003) Tests and tolerances for high-performance software-implemented fault detection. *IEEE Transactions on Computers* 52(5): 579–591.
- Van Horn J (2005) Towards achieving relentless reliability gains in a server marketplace of teraflops, laptops, kilowatts, and “cost, cost, cost”...: Making peace between a black art and the bottom line. In: *Proceedings of the IEEE international test conference (ITC)*, p. 8.
- Wang N and Patel S (2006) ReStore: Symptom-based soft error detection in microprocessors. *IEEE Transactions on Dependable and Secure Computing* 3(3): 188–201.
- Wittgenstein L (1953) *Philosophical Investigations.*: The Macmillan Company, New York.
- Yang J, Zhang M, Strachan J, et al. (2010) High switching endurance in TaO_x memristive devices. *Applied Physics Letters* 97(23): 232102.
- Young JW (1974) A first order approximation to the optimum checkpoint interval. *Communications of the ACM* 17(9): 530–531.
- Yu J, Garzaran MJ and Snir M (2009) Esoftcheck: Removal of non-vital checks for fault tolerance. In: *Proceedings of the 7th annual IEEE/ACM international symposium on code generation and optimization*, pp. 35–46.
- Yu S, Yin Chen Y, Guan X, et al. (2012) A Monte Carlo study of the low resistance state retention of HfO_x based resistive switching memory. *Applied Physics Letters* 100(4): 043507.
- Zhang M, Mitra S, Mak TM, et al. (2006) Sequential element design with built-in soft error resilience. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 14(13): 1368–1378.
- Zheng G, Ni X and Kalé L (2012) A scalable double in-memory checkpoint and restart scheme towards exascale. In: *Proceedings of the IEEE/IFIP international conference on dependable systems and networks (DSN)*, pp. 1–6.
- Zhou J, Wang M and Wong M (2010) Instability of p-channel poly-Si thin-film transistors under dynamic negative bias temperature stress. In: *17th IEEE international symposium on the physical and failure analysis of integrated circuits (IPFA)*, pp. 1–4.
- Zio E, Maio FD and Stasi M (2010) A data-driven approach for predicting failure scenarios in nuclear systems. *Annals of Nuclear Energy* 37: 482–491.

Author biographies

Marc Snir is the Director of Argonne’s Mathematics and Computer Science Division and the Michael Faiman and Saburo Muroga Professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign. His research is focused on HPC, with recent work on programming models, performance analysis, and resilience. Snir received his PhD from the Hebrew University of Jerusalem. He spent time at NYU, where he worked on the NYU Ultracomputer, and at IBM Research, where he led the research team that worked on the software for the IBM SP and Blue Gene systems. At UIUC, he headed the CS department and led the creation of the Illinois Informatics Institute. Marc Snir is an AAAS, ACM, IEEE, and Argonne Fellow. He has recently received the IEEE Award for Excellence in Scalable Computing and the IEEE Computer Society Seymour Cray Computer Engineering Award.

Pavan Balaji holds appointments as a Computer Scientist at the Argonne National Laboratory, as an Institute Fellow of the Northwestern-Argonne Institute of Science and Engineering at Northwestern University, and as a Research Fellow of the Computation Institute at the University of Chicago. He leads the Programming Models and Runtime Systems group at Argonne. His research interests include parallel programming models and runtime systems for communication and I/O, modern system architecture (multicore, accelerators, complex memory subsystems, high-speed networks), and cloud computing systems. He has nearly 100 publications in these areas and has delivered nearly 120 talks and tutorials at various conferences and research institutes. He is a recipient of several awards including the U.S. Department of Energy Early Career award in 2012, TEDx Midwest Emerging Leader award in 2013, Crain’s Chicago 40 under 40 award in 2012, Los Alamos National Laboratory Director’s Technical Achievement award in 2005, Ohio State University Outstanding Researcher award in 2005, five best-paper awards, and various others. He serves as the worldwide chairperson for the IEEE Technical Committee on Scalable Computing (TCSC). He has also served as a chair or editor for nearly 50 journals, conferences, and workshops, and as a technical program committee member in numerous conferences and workshops. He is a senior member of the IEEE and a professional member of the ACM. More details are available at <http://www.mcs.anl.gov/~balaji>.

Todd Munson received a BS in Computer Science from the University of Nebraska in 1995, and an MS in 1996 and PhD in 2000 in Computer Science from the University of Wisconsin at Madison. He is a Computational Scientist in the Mathematics and Computer Science Division at Argonne National Laboratory, a Senior Fellow in the Computation Institute at the University of Chicago and Argonne National Laboratory. The primary focus of his research is algorithms and applications of numerical optimization and variational inequalities. He has been widely recognized for his contributions. Among other honors he was awarded a Presidential Early Career Award for Scientists and Engineers from the White House, an Early Career Scientist and Engineer Award from the U.S. Department of Energy in 2006, and the Beale-Orchard-Hayes Prize from the Mathematical Programming Society in 2003. He has twice been invited to the White House to meet the President of the United States (Bush 41 and Bush 43).

Andrew A Chien is the William Eckhardt Professor in Computer Science at the University of Chicago. He is also a Senior Fellow at UC's Computation Institute and a Senior Computer Scientist at Argonne National Laboratory. His research interests include parallel computing, computer architecture, and cloud computing. From 2005 to 2010, Chien was Vice President of Research at Intel Corporation where he launched new initiatives in parallel software, mobile computing, cloud computing, and exascale research. From 1998 to 2005, Chien was the SAIC Endowed Chair Professor in the Department of Computer Science and Engineering where he founded the Center for Networked Systems at the University of California San Diego. From 1990 to 1998, he was a Professor of Computer Science at the University of Illinois at Urbana-Champaign and the National Center for Supercomputing Applications (NCSA). He has served on numerous advisory committees for the National Science Foundation, Department of Energy, and universities such as Stanford, EPFL, and Cal-Berkeley. Chien earned BS, MS, and PhD degrees at the Massachusetts Institute of Technology, and is a Fellow of the ACM, IEEE, and AAAS.

Pradip Bose is a research scientist at IBM T. J. Watson Research Center, where he manages a department on power-efficient, resilient systems. He holds a PhD from the University of Illinois at Urbana-Champaign. He has been associated with the definition and pre-silicon modeling of virtually all POWER-series processors, beginning with the original pre-product super scalar RISC project at IBM. He is a member of IBM's Academy of Technology and an IEEE Fellow.

Al Geist is a Corporate Research Fellow at Oak Ridge National Laboratory. He is the Chief Technology

Officer of the Leadership Computing Facility. His recent research is on exascale computing and resilience needs of hardware and software.

Saurabh Bagchi is a Professor in the School of Electrical and Computer Engineering and the Department of Computer Science (by courtesy) at Purdue University in West Lafayette, Indiana. He is a Senior Member of IEEE and ACM, a Distinguished Speaker for ACM, an IMPACT Faculty Fellow at Purdue (2013–14), and an Assistant Director of the CERIAS security center at Purdue. He leads the Dependable Computing Systems Laboratory (DCSL), where his group performs research in practical system design and implementation of dependable distributed systems. Since 2011, he has been serving as a Visiting Scientist with IBM Austin Research Lab.

Mattan Erez is an Associate Professor at the Department of Electrical and Computer Engineering at the University of Texas at Austin. His research focuses on improving the performance, efficiency, and scalability of computing systems through advances in hardware architecture, software systems, and programming models. The vision is to increase the cooperation across system layers and develop flexible and adaptive mechanisms for proportional resource usage. Erez received a BSc in Electrical Engineering and a BA in Physics from the Technion, Israel Institute of Technology, and his MS and PhD in Electrical Engineering from Stanford University.

Sarita V Adve is Professor in Computer Science at the University of Illinois. Her research interests are broadly in computer architecture and systems. She leads the SWAT project, one of the early projects to explore holistic software-driven solutions for hardware resiliency. She is an ACM Fellow, an IEEE Fellow, and an ABI Women of Vision award winner in innovation.

Sven Leyffer is a senior computational mathematician in the Mathematics and Computer Science Division at Argonne National Laboratory, and a Senior Fellow of the Computation Institute. He obtained his PhD from the University of Dundee, UK, and has held postdoc positions at Dundee, Northwestern University, and Argonne. He is a Fellow of the Society for Industrial and Applied Mathematics.

Nathan DeBardleben received his PhD in Computer Engineering from Clemson University in 2004 and started at Los Alamos National Laboratory the same year. DeBardleben has been influential in defining the field of HPC resilience, its challenges and potentials. He has co-authored a handful of governmental position papers on the subject as well as his own research publications. In his own research, his focus is on characterizing the impact of soft errors on systems and

applications. DeBardeleben is on numerous reliability program committees, runs his own workshop (Fault-tolerance for HPC at Extreme Scale (FTXS)) and runs the Los Alamos National Laboratory resilience site (<http://institute.lanl.gov/resilience/>).

Christian Engelmann is Task Lead of the System Software Team in the Computer Science and Mathematics Division at Oak Ridge National Laboratory. He earned his PhD in Computer Science in 2008 and his MSc in Computer Science in 2001, both from the University of Reading, UK. He also obtained a German Certified Engineer diploma in Computer Systems Engineering in 2001 from the University of Applied Sciences, Berlin. Engelmann's research aims at computer science challenges for extreme-scale HPC system software, such as dependability, scalability, and portability. His primary expertise is in HPC resilience, that is, providing efficiency and correctness in the presence of faults, errors, and failures through avoidance, masking, and recovery. His secondary expertise is in HPC hardware/software co-design through light-weight simulation of extreme-scale systems with millions of processor cores to study the impact of hardware properties on parallel application performance.

Jim Belak is a senior scientist in the Condensed Matter and Materials Division at Lawrence Livermore National Laboratory. He is Co-PI and Deputy Director for the Exascale Co-design Center for Materials in Extreme Conditions (ExMatEx), a joint project with Los Alamos National Laboratory, ORNL, SNL-A, Stanford, and CalTech, funded by the DOE Office of Advanced Scientific Computing Research. The goal of ExMatEx is to use the supercomputer codes used to study matter under extreme conditions to guide the design of future supercomputers and use the understanding gained to refactor and create new supercomputer codes. He earned his PhD in Condensed Matter Physics from Colorado State University.

Fred Johnson is currently with SAIC serving as senior SAIC technical advisor to the DOE NNSA Advanced Simulation & Computing organization. He has retired as the Senior Technical Manager for Computer Science in DOE/ASCR where he was the Program Manager responsible for fundamental computer science research and research on high-performance system software and tools including programming models, debugging and performance evaluation tools, software component architectures for high-performance systems, and next-generation runtime and OSs.

Pedro Diniz earned his PhD from the University of California, Santa Barbara, in Computer Science in 1997. Since then he has been a Research Assistant Professor of Computer Science with the University of Southern California in Los Angeles, California. He has

also been involved in several research projects focusing on programming technology and execution models addressing productivity-related issues as well as fault-tolerance for large-scale high-performance architectures. He has participated in various scientific proposal review boards at the National Science Foundation as well as at the European Commission in Brussels. Over the last 20 years he has been heavily involved in the scientific community having participated as part of the technical program committee of over 15 international conferences in the area of HPC, reconfigurable and field-programmable computing.

Paul Coteus is an IBM Fellow in the Systems Department at the Thomas J. Watson Research Center. Coteus completed his PhD in Physics at Columbia University and joined IBM in 1988, leaving his position as Assistant Professor of Physics at the University of Colorado. He has directed and designed advanced packaging for high-speed electronics, memory systems, and processor complexes. He is currently the Chief Engineer of Data Centric Systems, and also leads the system engineering for the full line of Blue Gene Supercomputers, honored in 2008 with the National Medal of Technology and Innovation. He is an IEEE Fellow, a member of IBM's Academy of Technology, and an IBM Master Inventor. He has authored more than 90 papers in the field of electronic packaging, and holds over 120 US patents.

Rinku Gupta is a senior scientific developer at Argonne National Laboratory. She received her MS degree in Computer Science from Ohio State University in 2002. She has several years of experience developing systems and infrastructure for enterprise HPC. Her research interests primarily lie towards middleware libraries, programming models, and designing fault-tolerance frameworks in HEC systems. More details about her are available at <http://www.mcs.anl.gov/~rgupta>.

Franck Cappello holds a Senior Computer Scientist position at Argonne National Laboratory where he leads the resilience effort. He is the main PI of the G8 'Enabling Climate Simulation at Extreme Scale' project gathering research groups from six countries. He is also the initiator and co-director of the INRIA-Illinois-ANL Joint Laboratory on Petascale Computing. Before moving to USA, he led the Grand-Large and Grid'5000 projects in France at INRIA, focusing on high-performance issues and research methodology for large-scale distributed systems. He has authored more than 130 papers and contributed to more than 70 program committees. He is an editorial board member of the international Journal of Grid Computing, Journal of Grid and Utility Computing, and Journal of Cluster Computing. He served in the steering committees of IEEE HPDC and

IEEE/ACM CCGRID. He is the Program Co-Chair of ACM HPDC 2014 and ACM CAC 2014.

Rob Schreiber is a Distinguished Technologist at Hewlett Packard Laboratories. Schreiber's research spans sequential and parallel algorithms for matrix computation, compiler optimization for parallel languages, and high-performance computer design. With Moler and Gilbert, he developed the sparse matrix extension of Matlab. He created the NAS CG parallel benchmark. He was a designer of the High Performance Fortran language. At HP, he led the development of PICO, a system for synthesis of custom hardware accelerators. His recent work concerns architectural uses of CMOS nanophotonic communication and NVM architecture. He is an ACM Fellow, a SIAM Fellow, and was awarded, in 2012, the Career Prize from the SIAM Activity Group in Supercomputing.

Dean Liberty is a Fellow at Advanced Micro Devices (AMD). He leads the Reliability/Availability/Serviceability (RAS) Architecture and Strategy team, focusing on long-term planning, detailed architecture, and short-term implementation for resilience in AMD processors. Dean has been in the computer industry for over 30 years, and involved in HPC systems for over 20 years. His experience covers a range of hardware and software, and his interests lie in bridging the gap between the two.

Eric Van Hensbergen is currently a principal design engineer at ARM Research in Austin, Texas. His current research focuses on exploring energy-efficient approaches to HPC through balance-driven co-design. Previous to ARM he was a research staff member in the Future Systems department at IBM's Austin Research Lab. Over 12 years at IBM Research, he worked on distributed OSs for HPC, low-power dense server and network processor appliance blades, DRAM power management, full system simulation, HPC, hypervisors, and the Linux OS. Before coming to IBM, he worked for four years at Lucent Technologies Bell Laboratories on the Plan 9 and Inferno OSs.

Sriram Krishnamoorthy received his BE degree from the College of Engineering-Guindy, Anna University, Chennai, and his MS and PhD degrees from The Ohio State University, Columbus, Ohio. He is currently a research scientist at Pacific Northwest National Laboratory. His research focuses on parallel programming models, fault tolerance, and compile-time/runtime optimizations for HPC. He has over 60 peer-reviewed conference and journal publications, receiving best-paper awards for his publications at the International Conference on High Performance Computing (HiPC'03) and the International Parallel and

Distributed Processing Symposium (IPDPS'04). He is a recipient of the U.S. Department of Energy Early Career award and Pacific Northwest National Laboratory's Ronald L. Brodzinski Award for Early Career Exceptional Achievement in 2013. He is a senior member of the IEEE and a professional member of ACM.

Subhasish Mitra directs the Robust Systems Group in the Department of Electrical Engineering and the Department of Computer Science of Stanford University, where he is the Chambers Faculty Scholar of Engineering. Before joining Stanford, he was a Principal Engineer at Intel Corporation. His research interests include robust system design, VLSI design, CAD, validation and test, and emerging nanotechnologies. His research results have seen widespread proliferation in industry, and have been recognized by several prestigious awards including the Presidential Early Career Award for Scientists and Engineers from the White House, the Intel Achievement Award, Intel's highest corporate honor, and several best-paper awards for publications at major conferences and journals. He is a Fellow of the IEEE.

Jon Stearley is a senior member of technical staff at Sandia National Laboratories. His interests include historical and live mining of system logs to identify the root causes of faults, the propagation of errors, and their effects on user jobs, towards faster fixes today and better designs tomorrow.

Saverio Fazzari works for Booz Allen acting as a senior technical advisor to DARPA and other government agencies for numerous programs. Fazzari has a strong background in all areas of semi-conductor design and fabrication, from algorithm development through device implementation. His specialty is advanced circuit design and development strategies with a focus on hardware cyber security issues including trusted design and fabrication. His experience includes extensive commercial experience, leading production innovation and development across all facets of the electronic design process.

Jacob A Abraham is a Professor in the Department of Electrical and Computer Engineering at the University of Texas at Austin. He is also director of the Computer Engineering Research Center and holds a Cockrell Family Regents Chair in Engineering. He received a bachelor degree in Electrical Engineering from the University of Kerala, India, in 1970. His MS degree, in Electrical Engineering, and PhD, in Electrical Engineering and Computer Science, were received from Stanford University, California, in 1971 and 1974, respectively. From 1975 to 1988 he was on the faculty of the University of Illinois, Urbana, Illinois.

William Carlson is a member of the research Computing Sciences Staff at the IDA Center for Computing Sciences where, since 1990, his focus has been on applications and system tools for large-scale parallel and distributed computers. He also leads the UPC Language Effort, a consortium of industry and academic research institutions aiming to produce a unified approach to parallel C programming based on global address space methods. Carlson graduated from Worcester Polytechnic Institute in 1981 with a BS degree in Electrical Engineering. He then attended Purdue University, receiving MSEE and PhD degrees in Electrical Engineering in 1983 and 1988, respectively. From 1988 to 1990, Carlson was an Assistant Professor at the University of Wisconsin–Madison, where his work centered on performance evaluation of advanced computer architectures.

Robert W Wisniewski is an ACM Distinguished Scientist and the Chief Software Architect for Extreme-Scale Computing and a Senior Principal Engineer at Intel Corporation. He has published over 60 papers in the area of HPC, computer systems, and system performance, and has filed over 50 patents. Before coming to Intel, he was the chief software architect for Blue Gene Research and manager of the Blue Gene and exascale research software team at the IBM T.J. Watson Research Facility, where he was an IBM Master Inventor and lead the software effort on Blue Gene/Q, which was the fastest machine in the world on the June 2012 Top 500 list, and occupied four of the top 10 positions. Prior to working on Blue Gene, he worked on the K42 scalable OS project targeted at scalable next-generation servers and the DARPA HPCS project on continuous program optimization that utilizes integrated performance data to automatically improve application and system performance. Before joining IBM Research, and after receiving a PhD in Computer Science from the University of Rochester, he worked at Silicon Graphics on high-end parallel OS development, parallel real-time systems, and real-time performance monitoring.

Appendix A Derivation of optimal checkpoint interval

We assume a global checkpointing model: the system is periodically taking global checkpoints; after a failure, computation is restarted from the last checkpoint. We use the following parameters:

- Checkpoint time is C .
- Recovery time is R .
- Checkpoint interval is τ : a new checkpoint is taken time τ after the previous checkpoint started, or time τ after a failure occurred.

- Probability of failure within a time interval τ is $F(\tau)$.
- Time to first failure *given that a failure occurs* within the interval τ is $W(\tau)$.

We assume that C , R , and τ are constant, while $W(\tau)$ is a random variable. We further assume that the system is *memoryless*: $F(\tau)$ and $W(\tau)$ are the same, for each time interval.

We divide the computation into *epochs*: a new epoch starts when a failure occurred, or when a checkpoint completed. Let $Comp_i$ be the amount of useful computation done in epoch i and let $Time_i$ be the amount of wallclock time consumed by epoch i . $Comp_i$ are i.i.d. random variables and $Time_i$ are i.i.d. random variables ($Comp_i$ and $Time_i$ are not independent).

We have

$$Comp = \begin{cases} \tau - C & \text{if epoch completes normally} \\ -R & \text{otherwise} \end{cases} \quad (1)$$

The $-R$ represents the fact that not only was no progress made, but the computation now requires recovery. Also,

$$Time = \begin{cases} \tau & \text{if epoch completes normally} \\ W(\tau) & \text{otherwise} \end{cases} \quad (2)$$

We define the *utilization* of the system to be ratio between compute time and wall-clock time:

$$Util = \lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n Comp_i}{\sum_{i=1}^n Time_i}$$

We have $(1/n) \sum_{i=1}^n Comp_i \rightarrow E[Comp]$ and $(1/n) \sum_{i=1}^n Time_i \rightarrow E[Time]$, so that

$$Util = \frac{E[Comp]}{E[Time]} \quad (3)$$

We derive, from equations (1) and (2),

$$E[Comp] = (1 - F(\tau))(\tau - C) - F(\tau)R$$

and

$$E[Time] = (1 - F(\tau))\tau + F(\tau)E[W(\tau)]$$

so that

$$Util = \frac{(1 - F(\tau))(\tau - C) - F(\tau)R}{(1 - F(\tau))\tau + F(\tau)E[W(\tau)]} \quad (4)$$

Note that this formula does not involve any approximation and does not depend on the distribution of between-failure intervals.

We shall assume from now on that failures occur according to a Poisson process, and normalize time so that MTTF equals 1 (i.e. we express checkpoint time and recovery time as fractions of MMTF). Thus,

$$F(\tau) = 1 - e^{-\tau}$$

and

$$E[W(\tau)] = \frac{1}{F(\tau)} \int_0^\tau x e^{-x} dx$$

But

$$\int_0^\tau x e^{-x} dx = -(x + 1)e^{-x/1} \Big|_0^\tau = -(\tau + 1)e^{-\tau} + 1$$

Thus

$$E[Comp] = e^{-\tau}(\tau - C) - (1 - e^{-\tau})R = (\tau - C + R)e^{-\tau} - R$$

$$E[Time] = \tau e^{-\tau} - (\tau + 1)e^{-\tau} + 1 = 1 - e^{-\tau}$$

and

$$Util = \frac{(\tau - C + R)e^{-\tau} - R}{1 - e^{-\tau}}$$

We want to select the τ that maximizes utilization. Such τ solves the equation

$$\frac{dUtil}{d\tau} = 0$$

We compute derivatives and obtain the equation

$$(e^{-\tau} - (\tau - C + R)e^{-\tau})(1 - e^{-\tau}) - ((\tau - C + R)e^{-\tau} - R)e^{-\tau} = 0$$

Simplifying, we obtain the equation

$$e^{-\tau} = 1 - \tau + C \tag{5}$$

and

$$Util = \frac{(\tau - C + R)(1 - \tau + C) - R}{\tau - C} = 1 - \tau + C - R \tag{6}$$

We solve equation (5) numerically, for different values of M and R , and plug into equation (6) in order to compute the best possible utilization, as a function of (relative) MTTF and recovery time.

Various approximations can be derived from equation (5): if we approximate e^x with the first three terms of its Taylor expansion, then we get

$$1 - \tau - \frac{\tau^2}{2} = 1 - \tau + C \tag{7}$$

so that $\tau_{opt} = \sqrt{2C}$ and $Util = 1 - \sqrt{2C} + C - R$.

If the MTBF is M (rather than 1), we get

$$\tau_{opt} = M\sqrt{2C/M} = \sqrt{2CM}$$

and

$$Util = 1 - \sqrt{2C/M} + \frac{C - R}{M}$$

The approximation is valid when $C \ll M$ (Young, 1974). Higher-level approximations are derived in Daly (2006).