

# ORNL REPORT

ORNL/TM-2013/283

Unlimited Release

Printed July 2015

## Numerical analysis of fixed point algorithms in the presence of hardware faults

Miroslav K. Stoyanov, Clayton G. Webster

Prepared by  
Department of Computational and Applied Mathematics  
Computer Science and Mathematics Division  
Oak Ridge National Laboratory  
One Bethel Valley Road, Oak Ridge, Tennessee 37831

The Oak Ridge National Laboratory is operated by UT-Battelle, LLC,  
for the United States Department of Energy under Contract DE-AC05-00OR22725.  
Approved for public release; further dissemination unlimited.

## DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via the U.S. Department of Energy (DOE) Information Bridge.

**Web site** <http://www.osti.gov/bridge>

Reports produced before January 1, 1996, may be purchased by members of the public from the following source.

National Technical Information Service

5285 Port Royal Road

Springfield, VA 22161

Oak Ridge, TN 37831

**Telephone** 703-605-6000 (1-800-553-6847)

**TDD** 703-487-4639

**Fax** 703-605-6900

**E-mail** [info@ntis.gov](mailto:info@ntis.gov)

**Web site** <http://www.ntis.gov/support/ordernowabout.htm>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange (ETDE) representatives, and International Nuclear Information System (INIS) representatives from the following source.

Office of Scientific and Technical Information

P.O. Box 62

Oak Ridge, TN 37831

**Telephone** 865-576-8401

**Fax** 865-576-5728

**E-mail** [reports@osti.gov](mailto:reports@osti.gov)

**Web site** <http://www.osti.gov/contact.html>

## NOTICE

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

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



# NUMERICAL ANALYSIS OF FIXED POINT ALGORITHMS IN THE PRESENCE OF HARDWARE FAULTS

M. Stoyanov<sup>\*</sup> C. Webster<sup>†</sup>

## Abstract.

The exponential growth of computational power of the extreme scale machines over the past few decades has led to a corresponding decrease in reliability and a sharp increase of the frequency of hardware faults. Our research focuses on the mathematical challenges presented by the silent hardware faults. i.e., faults that can perturb the result of computations in an inconspicuous way. Using the approach of selective reliability we present an analytic fault mode that can be used to study the resilience properties of a numerical algorithm. We apply our approach to the classical fixed point iteration and demonstrate that in the presence of hardware faults, the classical method fails to converge in expectation. We present a modified resilient algorithm that detects and rejects faults resulting in error with large magnitude, while small faults are negated by the natural self-correcting properties of the algorithm. We show that our method is convergent (in first and second statistical moments) even in the presence of silent hardware faults.

**Key words.** algorithmic resilience, fixed point method, linear solvers, numerical linear algebra, fixed point methods

**1. Introduction.** As modern science and engineering push the limits of our knowledge and understanding, there is an ever increasing demand for more complex mathematical models to describe various real world phenomena. Coupling between processes and scales as well as the need for better accuracy have created an insatiable demand for increased computing power. To this end, bigger and faster computers are proposed and built every year and every one of those machines is designed to challenge the limits of our configurations. Transistor's fabrication size shrinks exponentially and constraints on the power consumption force circuits to operate at near critical voltage, so that even a small power fluctuation may have an adverse effect on operation. In addition, the ever increasing number of components offers more and more opportunity for hardware failure. Unfortunately, the rapid increase of computational power has led to a corresponding decrease in reliability, which has created an ever growing concern for the reliability of the computed results [5, 7, 14, 24–26].

In this work we focus on the problem of *silent* hardware faults that can perturb the result of computations in an inconspicuous way. We devise a theoretical fault model and a framework for analysis of error propagation, as well as assessment and enhancement of algorithm's resilience. In what follows, we describe the challenges associated with hardware failure and silent faults, that have led to the research described in this effort.

**1.1. Hardware failure.** Hardware faults are a well known issue and measures are taken to ensure that every computer component operates reliably. For example, memory comes with error detection and correction [22, 26] and Cyclic-Redundancy-Checksums (CRC) are used to guard against data corruption in network traffic. However, processing cores remain largely unprotected [14, 17] and adding redundancy to all chips in a cluster is very expensive in terms of manufacturing cost and increased energy consumption. Faults in a single

---

<sup>\*</sup>Department of Computational and Applied Mathematics, Oak Ridge National Laboratory, One Bethel Valley Road, P.O. Box 2008, MS-6367, Oak Ridge, TN 37831-6164 ([stoyanovmk@ornl.gov](mailto:stoyanovmk@ornl.gov)).

<sup>†</sup>Department of Computational and Applied Mathematics, Oak Ridge National Laboratory, One Bethel Valley Road, P.O. Box 2008, MS-6164, Oak Ridge, TN 37831-6164 ([webstercg@ornl.gov](mailto:webstercg@ornl.gov))

component are rare, however, in an extreme-scale machine, small chances of failure are multiplied so they cannot be ignored. Experiments performed in 2011 on Jaguar (the largest supercomputer at the time) showed a single node failure almost daily [14] and hundreds of (corrected) bit-flips every hour [25]. Hardware level error detection and correction are indispensable tools, however, the number of undetected and uncorrected faults is constantly increasing.

Software based resilience techniques are used to complement the hardware approach. One of the most common types of hardware fault is the fail-stop (sometimes called *hard* faults), where either the entire application, or some of the nodes of the cluster, crash and fail to return a result. This problem has been the focus of research in computational science over the past decades and the most common approach for addressing this issue is the checkpoint restart methods [2, 13]. Practically, all modern code is written so that it periodically saves the state of the software and in the event of a crash, computations can be resumed from the last saved checkpoint.

**1.2. Silent faults.** Hardware malfunction that doesn't result in a crash, may result in data corruption. If the fault affects a permanent data structure (e.g., an iteration matrix), then checksums can be used to detect the deviation, however, not all faults leave a "visible" permanent trace. For example, suppose two variables with values 2 (binary 010) and 4 (binary 100) need to be added. First they are moved from the DRAM into the SRAM (CPU cache). Then suppose the hardware suffers a voltage fluctuation causing the cache to fail and flip one bit in the second variable making it into a 6 (binary 110). When the add instruction is executed, it will result in the incorrect answer 8 (binary 1000), which will be stored back into DRAM and the corrupt data in the cache will be cleared. Effectively, the hardware computes erroneously that  $2 + 4 = 8$ . If the corrupted value contains invalid data (e.g., *inf* or *nan*) or if the values fall outside of physical boundaries (e.g., negative area), then some indication of a fault is still present. However, if the result is not "obviously" wrong, then we have the most sinister type of a fault that perturb the computations in way that is incorrect but yet believable. In our work, we call those the *silent* faults.

The main focus of this work is to propose and analyze novel numerical methods that converge even in the presence of silent faults, regardless of their source. We do not restrict our attention to bit-flips and our approach extends to any type of perturbation of the numerical computations that does not permanently alter the persistent data in main memory. The name *silent* is appropriate here, since it emphasizes the main challenge associated with detection of those faults.

Experimental studies of the effects of faults (silent or otherwise) have been performed (e.g., [6, 9]) and techniques for fault detection have been proposed (e.g., [4, 19]). Simple enhancements can improve the overall resilience of the methods, however, not all silent faults can be corrected or even detected, and experiments provide insight on the reliability but are restricted to the specific implementations and specific faults tested (e.g., bit-flips). In the context of linear algebra, Algorithm Based Fault Tolerance (ABFT) [1, 16, 18] is an advanced and very widely utilized technique. ABFT uses check-sums to detect data corruption in matrices and vectors as well as correct some silent faults in the basic algebraic operations, and ABFT has been successfully applied to dense [8] and sparse [23] matrix operations on large scale systems. However, the application of ABFT remains restricted to linear opera-

tions with explicit matrices, specific data-structures and implementation. Checksumming is a very powerful way to guard against data corruption in persistent structures, however, a more general higher-level approach is needed to handle silent faults in the broader class of problems.

The work [3, 15] propose an alternative approach where the program execution is split into two modes, denoted *safe* and *fast* (respectively *host* and *guest* in [15]). The computations in *safe* mode are assumed to be performed reliably (i.e., zero or negligible probability of a fault), however, they are also more expensive. The *safe* mode has to be implemented with either redundant computation or via the use of specialized hardware, hence only the minimum amount of work should be performed in this mode. On the other hand, the *fast* mode is cheap, but prone to hardware failure. This approach is applied to a modified version of the restarted GMRES algorithm, where a convergence test is performed in *safe* mode between every two outer GMRES iterations [15]. The convergence test rejects the next iterate unless the residual error is reduced, thus large errors are rejected and only small ones are accepted by the algorithm. Similar approach has also been applied to Conjugate-Gradient method [21].

The approach of selective reliability is a viable technique for dealing with silent faults, however, it also presents a number of new challenges. For example, the convergence of the algorithm deteriorates as the frequency of the faults increases, due to the additional work needed to correct each fault. However, the rate of deterioration, and the relationship between the properties of the matrix (e.g., condition number), and the type of errors that can be resolved remains unknown. Thus, the frequency and type of faults that can be resisted is not fully characterized. Even if convergence is achieved, in the presence of random faults, the final result of the computations is a random variable with associated probability distribution, and the statistics of the distribution need to be quantified. As such, in what follows, a novel numerical analysis paradigm is presented in which the potential of selective reliability can be rigorously analyzed.

**1.3. Analysis framework in the presence of faults.** Our modern treatment of predicting the behavior of physical phenomena and complex engineering problems, relies on mathematical modeling followed by computer simulations. Traditional numerical algorithms have been developed with the assumption that the underlying algebraic operations, e.g., matrix vector products, can be carried out accurately. However, as the reliability of supercomputers decrease, scientists and engineers can no longer afford to make this assumption; especially when providing analysis and informing decision makers of the behavior of computer simulations. Numerical algorithms compute approximations to the solutions of mathematical models, and therefore, introduce an associated numerical error. Silent faults introduce additional hardware artifacts in the computer simulation that are often times indistinguishable from the numerical error, and therefore, a new analysis paradigm to simultaneously consider both numerical and hardware error must be established.

In this work, we establish a framework of numerical analysis that is free from the assumption that computations can always be performed accurately. Following the approach of selective reliability [3, 15], we assume that the result of numerical operations will be randomly perturbed. In order to provide analytic framework that is as general as possible to ever changing hardware architectures and implementation techniques, we make minimal as-

sumptions on the structure of such perturbations. In the new analysis paradigm, we provide rigorous error bounds with respect to the hardware error, and we prove convergence of the algorithms, even if they are implemented on unreliable hardware (i.e., hardware susceptible to silent faults). As a demonstration of our approach on solutions of systems of equations, we analyze the widely used fixed point family of iterative solvers. We show how error of large magnitude can be rejected, so that faults can only introduce discrepancy that is of the order of the standard numerical error, thus preserving the theoretical convergence rates. This is the first step towards building a suite of resilient numerical algorithms that provide reliable results even in an unreliable computing environment.

The rest of the paper is organized as follows: in §2 we present the fault model and define convergence with respect to hardware induced error. In §3, we present the analysis of the classical fixed point iteration and how it can be enhanced using the technique of selective reliability. In §4, we present a numerical comparison between the classical and enhanced fixed point methods.

**2. Hardware error propagation and convergence.** In this section we establish the framework for analyzing the effects of silent hardware faults on numerical algorithms. We propose a fault model that is agnostic to specific hardware architectures or software implementations. Using this model, we provide rigorous definition of *hardware error* and convergence in the presence of silent hardware malfunction.

**2.1. Fault model.** Hardware faults can be very diverse in nature and the resulting discrepancy introduced into the computations can be unpredictable. For example, consider some of the possible effects of a bit-flip, which is one of the most common types of hardware faults. A bit-flip in the opcode may convert one numerical operation into another (i.e., addition may turn into multiplication). A bit-flip in a floating point number can cause error in the entire range of values representable by the floating point standard, e.g., see our earlier work on the IEEE-754 double precision representation [11, 12]. A bit-flip in a pointer or array index can cause data to be read from a random location in memory and hence result in an error of unpredictable pattern. In complex operations, such as sparse matrix vector product, a fault in one number can propagate to multiple entries of the resulting vector depending on the sparsity pattern. Furthermore, some of the most efficient checkpoint/restart techniques for handling hard faults use partial/incomplete data recovery, that converts hard faults to soft faults and introduces additional error. Thus, the statistical properties of this new error strongly depend on the hardware circuits, floating point representation, data structures used, problem structure (e.g., matrix pattern) and even random memory content.

In order to keep our analysis as general as possible, we present a fault model that is agnostic with respect to all of the aforementioned specific conditions. First, consider that algorithms are comprised of a series of steps, where each step computes an intermediate result that leads to a final answer. Then, we assume that each algorithm is implemented with selective reliability, i.e., the steps are performed in either *safe* or *fast* modes, where the possibility of failure in safe mode is negligible, and every operation performed in fast mode is susceptible to hardware malfunction. The implementation of safe mode can be achieved via either redundancy or dedicated specialized hardware. Therefore, we assume that safe computations are significantly more expensive, and in this effort we require that the bulk

of the work is performed in fast mode.

We model the hardware faults and the associated random perturbations as a two level probability model, i.e., a mixture model. We assume that the hardware faults are independent random events, and for every step performed in fast mode there is a positive probability to encounter a hardware malfunction that perturbs the intermediate result. Thus, the first level is modeled by a Bernoulli discrete distribution, where the two options are “faults were encountered at this step” and “no faults were encountered at this step”. Once a fault is encountered, the magnitude and structure of the introduced discrepancy depends on too many specific factors, hence, we assume that the perturbation is a random variable with unknown probability density. We illustrate with an example below.

**Example 2.1** (Adding two vectors). *Consider the following three step algorithm:*

1. *Fast: Compute some  $\mathbf{x} \in \mathbb{R}^N$*
2. *Fast: Compute some  $\mathbf{y} \in \mathbb{R}^N$*
3. *Safe: Return  $\mathbf{z} = \mathbf{x} + \mathbf{y}$*

*The first two steps are performed in fast mode, while the last step is executed in safe mode. For steps 1 and 2, let  $p_1, p_2 \in \mathbb{R}$  so that  $0 < p_1, p_2 < 1$ , be the Bernoulli parameters that indicate the corresponding probability of encountering hardware faults. Each fault rate depends on the specific hardware configuration, code implementation, as well as the computational complexity of the associated step, i.e., operations that are computationally more expensive take longer to execute and are more likely to encounter faults. However, for a given algorithm and implementation on specific hardware, we assume that  $p_1$  and  $p_2$  are known, i.e., we assume that we can measure an approximate rate of faults of a given machine.*

*If the algorithm encounters a fault on the first step, instead of returning  $\mathbf{x}$  we would compute  $\mathbf{x} + \tilde{\mathbf{x}}$ , where  $\tilde{\mathbf{x}} \in \mathbb{R}^N$  is a random vector with unknown probability density. That is,  $\tilde{\mathbf{x}}$  is associated with a complete probability space  $(\mathbb{R}^N, \mathcal{F}, \tilde{P}_1)$ , where  $\mathcal{F} = 2^{\mathbb{R}^N}$  is the  $\sigma$ -algebra of events and  $\tilde{P}_1$  is some unknown probability distribution. We do not assume that  $\tilde{\mathbf{x}}$  has any specific structure or is bounded in any norm. Similarly, a hardware fault on the second step of the algorithm would perturb the result  $\mathbf{y}$  to  $\mathbf{y} + \tilde{\mathbf{y}}$ , where  $\tilde{\mathbf{y}}$  is associated with a complete probability space  $(\mathbb{R}^N, \mathcal{F}, \tilde{P}_2)$  and unknown probability measure  $\tilde{P}_2$ . In other words, the distribution of  $\tilde{\mathbf{x}}$  is a mixture between the distribution  $\tilde{P}_1$  with weight  $p_1$  and the Dirac-delta distribution  $\delta_0(\tilde{\mathbf{x}})$  with weight  $(1 - p_1)$ .*

*The last step of the algorithm is performed in safe mode and hence we assume that computations are done according to machine specification. However, the final answer  $\mathbf{z}$  would be affected by perturbations in  $\mathbf{x}$  and  $\mathbf{y}$ , and thus the result could be perturbed by either  $\tilde{\mathbf{x}}$  or  $\tilde{\mathbf{y}}$  or both  $\tilde{\mathbf{x}} + \tilde{\mathbf{y}}$ . Even if all the steps of the algorithm are deterministic, the stochastic nature of hardware faults turns the final result into a random variable.*

**2.2. Convergence with respect to hardware error.** The perturbation associated with a hardware failure in fast mode is distinct from round-off error. Finite precision standards such as IEEE-754 introduce rounding error with each operation, however, this error is small and deterministic in nature. Algorithm conditioning with respect to round-off error is a well established subject in numerical analysis, however, conditioning cannot account for

the random and potentially unbounded discrepancy introduced by hardware malfunction. Round-off error is introduced into the computations in compliance with standards and specifications, the hardware error that we consider is associated with abnormal machine behavior.

The error associated with numerical algorithms is usually measured in some norm. For all examples given in this paper, w.l.o.g., we assume that we are working with  $l^2$  norm defined on  $\mathbb{R}^N$ . However, our definitions are general and norm agnostic, and many of our results trivially extend to other norms. Thus, we present the following definition:

**Definition 1 (Hardware error).** *Error introduced into numerical computations by hardware malfunction.*

In Example 2.1, the perturbation associated with the first step is  $\tilde{\mathbf{x}}$  and, hence, the hardware error is  $\|\tilde{\mathbf{x}}\|$ . Similarly, the error associated with the second step is  $\|\tilde{\mathbf{y}}\|$  and the final hardware error is  $\|\tilde{\mathbf{x}} + \tilde{\mathbf{y}}\|$ . The random vectors  $\tilde{\mathbf{x}}$  and  $\tilde{\mathbf{y}}$  have unknown probability density, however, using our assumption about the probability of failure we have that  $\mathbb{P}_B(\tilde{\mathbf{x}} \neq 0) = p_1$  and  $\mathbb{P}_B(\tilde{\mathbf{y}} \neq 0) = p_2$ , where  $\mathbb{P}_B$  indicates the probability with respect to the Bernoulli measure. Therefore, the probability of non-zero final error is

$$\mathbb{P}_B(\tilde{\mathbf{x}} + \tilde{\mathbf{y}} \neq 0) = \mathbb{P}_B(\tilde{\mathbf{x}} \neq 0) + \mathbb{P}_B(\tilde{\mathbf{y}} \neq 0) - \mathbb{P}_B((\tilde{\mathbf{x}} \neq 0) \cap (\tilde{\mathbf{y}} \neq 0)) \leq p_1 + p_2 - p_1 p_2.$$

Note that we cannot make more specific conclusions about the distribution of  $\|\tilde{\mathbf{x}} + \tilde{\mathbf{y}}\|$  since we made no assumptions about the measures  $\tilde{P}_1$  and  $\tilde{P}_2$ .

In addition to hardware error, the accuracy of numerical algorithms is of critical importance. For example, a finite representation of an infinite dimensional operator introduces discretization error and the scheme is said to *converge* if the error goes to zero as the number of degrees of freedom increases. Another example is an iterative solver that generates a sequence of solutions. In this case, *convergence* is characterized by the sequence's ability to achieve a result as the number of iterations increases. Similarly round-off error can be reduced by using higher precision and representing real numbers with more bits, i.e., single, double, and quad precision.

In general, numerical convergence is related to the concept of additional computational works leading to an increasingly more accurate approximation. We want to extend this idea beyond numerical error and include hardware error, so that silent hardware faults are automatically corrected, accounting for additional computational complexity. The cost may be associated with extra iterations of an iterative solver or more refined mesh discretization, but in all cases more work should imply better accuracy. Furthermore, the hardware error is a random variable and hence, we consider accuracy in statistical sense, i.e., we want the statistical moments of the hardware error to tend to zero.

We formalize the idea of convergence in the presence of hardware faults by considering an algorithm with steps implemented in fast mode. We describe the hardware faults via the sequence of known Bernoulli parameters  $\mathbf{p} = (p_1, p_2, \dots, p_n, \dots)$ , and collection of unknown measures  $\tilde{\mathbf{P}} = (\tilde{P}_1, \tilde{P}_2, \dots, \tilde{P}_n, \dots)$ . Let  $e$  be the difference between the result of one realization of the algorithm and the desired exact solution, that is,  $e$  indicates the sum of numerical and hardware errors. Finally, denote by  $\mathbb{E}_{\mathbf{p}, \tilde{\mathbf{P}}}[e]$  and  $\mathbb{V}_{\mathbf{p}, \tilde{\mathbf{P}}}[e]$  the expectation and variance of  $e$  with respect to the mixture distributions defined by  $\mathbf{p}$  and  $\tilde{\mathbf{P}}$ . The statistical moments of  $e$  strongly depend on  $\tilde{\mathbf{P}}$ , however, in order to keep our approach agnostic to

specifics of the hardware, we take a conservative approach where we assume that every fault results in a the perturbation that would have the most detrimental effect on  $e$ . Following that notation, we define:

**Definition 2** (Convergence with respect to hardware error). *A method is convergent with respect to the hardware error, if for every  $\epsilon > 0$  there is a finite amount of work associated with the algorithm that will bound the statistics as*

$$\sup_{\tilde{\mathbf{P}}} \mathbb{E}_{\mathbf{p}, \tilde{\mathbf{P}}}[e] < \epsilon, \quad \text{and} \quad \sup_{\tilde{\mathbf{P}}} \mathbb{V}_{\mathbf{p}, \tilde{\mathbf{P}}}[e] < \epsilon^2, \quad (2.1)$$

where the supremum is taken over all possible probability measures on  $\mathbb{R}^N$ .

Traditional numerical analysis assumes that the hardware is perfectly reliable, i.e., all steps of algorithms are performed in safe mode and, therefore, the hardware error is always zero, i.e.,  $p = 0$  at all times. However, since the hardware error is potentially unbounded (i.e., supremum over  $\tilde{\mathbf{P}}$ ), even if a single step of an algorithm is performed in *fast* mode (i.e.,  $p > 0$ ), classical algorithms would fail to converge. For example, the simple approach described in Example 2.1 does not offer the ability to correct hardware error in the first two steps and hence the final error  $\tilde{\mathbf{x}} + \tilde{\mathbf{y}}$  will have non-zero expectation. To correct this and guarantee convergence, we must modify the first two steps in a way that allows us to increase the computational cost and make the statistical moments of  $\|\tilde{\mathbf{x}}\|$  and  $\|\tilde{\mathbf{y}}\|$  decay to zero.

Our objective is to take existing algorithms and allow for as many steps as possible in fast mode while still preserving convergence. In order to achieve this goal, we modify the algorithms in such a way as to control the propagation of faults so that the final hardware error has bounded statistics.

**3. Analytic framework for fixed point methods.** Iterative linear and non-linear solvers are the most widely utilized approaches for numerical solutions of large scale problems associated with most modern applications. Such methods generate a sequence of approximations of increasing fidelity to the exact solution, terminating after a convergence criteria has been achieved. One of the most widely used class of iterative methods is the family of fixed point solvers, which includes the linear methods such as Jacobi and Gauss-Siedel as well as nonlinear methods, such as Newton. Both types of methods possess properties that enable natural resiliency, and hence, are good example to illustrate our approach.

The basic idea of fixed-point methods is to construct a sequence of vectors  $\{\mathbf{x}^{(k)}\}_{k=0}^{\infty}$  that enjoy the property of convergence  $\mathbf{x}^{(k)} \rightarrow \mathbf{x}_G$ , where  $\mathbf{x}_G$  is a fixed point that satisfies the following linear or nonlinear system  $\mathbf{G}(\mathbf{x}_G) = \mathbf{x}_G$ . In practice, the fixed point iterative process is stopped at the minimum value  $n$ , such that  $\|\mathbf{x}^{(n)} - \mathbf{x}\| < \epsilon$ , where  $\epsilon$  is a fixed tolerance and  $\|\cdot\|$  is any convenient vector norm. However, since the exact solution is obviously not available, it is necessary to introduce a suitable stopping criteria. The increment  $e^{(k)} = \|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|$  between two successive iterates is the most common indicator for convergence, when the iteration reaches the asymptotic regime. Yet, this condition may cause premature termination in the first few iterations, therefore, we require that at least  $k_{min}$  number of iterations are taken<sup>1</sup>. We summarize the fixed point iteration in the following algorithm:

<sup>1</sup>Depending on the structure of  $\mathbf{G}(\mathbf{x})$ , in many applications it is sufficient to take  $k_{min} = 1$ .

**Algorithm 1** (Fixed point iteration).

Given  $\mathbf{G}$ ,  $\mathbf{b}$ , initial  $\mathbf{x}^{(0)}$ , tolerance  $\epsilon$  and  $k_{\min}$

Let  $k = 0$

**repeat**

$$\mathbf{x}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)})$$

$$e^{(k)} = \|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|$$

$$k = k + 1$$

**until** ( $k > k_{\min}$ ) **and** ( $e^{(k)} < \epsilon$ )

Convergence of Algorithm 1 is contingent upon the properties of  $G(\mathbf{x})$  and the initial iterate  $\mathbf{x}^{(0)}$ . One sufficient condition for convergence is based on the famous fixed point Theorem, a proof of which can be found in [20].

**Theorem 1** (Fixed point theorem). *Suppose  $\Gamma \subset \mathbb{R}^N$  is complete in some norm  $\|\cdot\|$  and let  $\mathbf{G} : \Gamma \rightarrow \Gamma$  be a contraction map, i.e., there is a constant  $r < 1$  so that*

$$\|\mathbf{G}(\mathbf{x}) - \mathbf{G}(\mathbf{y})\| \leq r\|\mathbf{x} - \mathbf{y}\|, \quad \text{for all } \mathbf{x}, \mathbf{y} \in \Gamma.$$

*Then, there is a unique fixed point  $\mathbf{x}_G \in \Gamma$  so that  $\mathbf{G}(\mathbf{x}_G) = \mathbf{x}_G$ . Furthermore, for any  $\mathbf{x}^{(0)} \in \Gamma$  the sequence  $\{\mathbf{x}^{(k)}\}_{k=0}^{\infty}$  defined by  $\mathbf{x}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)})$  converges to  $\mathbf{x}_G$  (i.e.,  $\|\mathbf{x}^{(k)} - \mathbf{x}_G\| \rightarrow 0$ ) and the error  $\|\mathbf{x}^{(k)} - \mathbf{x}_G\|$  is bounded by*

$$\|\mathbf{x}^{(k)} - \mathbf{x}_G\| \leq r^k \|\mathbf{x}^{(0)} - \mathbf{x}_G\|.$$

Theorem 1 is a strong result that not only guarantees convergence to the desired solution, but also gives the linear  $r^k$  decay of the error. Some fixed point methods exhibit even faster convergence rate, e.g., Newton's method converges quadratically, however, in this work, we focus on methods with linear rate.

**Example 3.1** (Linear solvers). *Here we focus on using fixed point methods to solve linear systems of type:*

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{3.1}$$

*where  $\mathbf{A} \in \mathbb{R}^{N \times N}$  is a given matrix,  $\mathbf{b} \in \mathbb{R}^N$  is a given right hand side vector and we are interested in the solution  $\mathbf{x}_G \in \mathbb{R}^N$  that satisfies (3.1). Two of the most common fixed point methods for linear systems are known as the Jacobi and Gauss-Siedel methods. The contraction maps and the corresponding contraction coefficients are given by:*

- **Jacobi Method**

$$\mathbf{G}_J(\mathbf{x}) = \mathbf{D}^{-1}(\mathbf{b} - \mathbf{S}\mathbf{x}), \quad r_J = \rho(\mathbf{D}^{-1}\mathbf{S}),$$

*where  $\mathbf{D} = \text{diag}(\mathbf{A})$  (i.e.,  $\mathbf{D}$  is a diagonal matrix with the diagonal entries of  $\mathbf{A}$ ),  $\mathbf{S} = \mathbf{A} - \mathbf{D}$ , and  $\rho$  indicates the spectral radius of the matrix  $\mathbf{D}^{-1}\mathbf{S}$ .*

- **Gauss-Seidel Method**

$$\mathbf{G}_{GS}(\mathbf{x}) = \mathbf{L}^{-1}(\mathbf{b} - \mathbf{U}\mathbf{x}), \quad r_{GS} = \rho(\mathbf{L}^{-1}\mathbf{U}),$$

*where  $\mathbf{L} = \text{lower}(\mathbf{A})$  (i.e.,  $\mathbf{L}$  is a lower triangular matrix of the entries of  $\mathbf{A}$ ),  $\mathbf{U} = \mathbf{A} - \mathbf{L}$ , and  $\rho$  indicates the spectral radius of the matrix  $\mathbf{L}^{-1}\mathbf{U}$ .*

The exact solution  $\mathbf{x}_G$  to (3.1) is a fixed point of both  $\mathbf{G}_J(\mathbf{x})$  and  $\mathbf{G}_{GS}(\mathbf{x})$ , therefore, if either  $r_J < 1$  or  $r_{GS} < 1$ , the corresponding fixed point iteration will converge to  $\mathbf{x}_G$  in  $l^2$  or any equivalent norm, for any  $\mathbf{x}^{(0)} \in \mathbb{R}^N$  (i.e.,  $\Gamma \equiv \mathbb{R}^N$ ).

**3.1. Hardware error analysis.** The most computationally expensive step of Algorithm 1 is finding the next iterate

$$\mathbf{x}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)}), \quad k = 0, 1, 2, \dots \quad (3.2)$$

In fact, in most applications this step is orders of magnitude more expensive than computing the increment  $e^{(k)}$ . However, hardware fault in the evaluation of  $e^{(k)}$  may cause premature termination of the algorithm, hence, we assume that the last two steps of the algorithm together with the conditional branching are computed in *safe* mode.

To reduce the computational cost associated with (3.2), we assume that the  $k + 1$  iterate is computed in fast mode. Then we must allow for the possibility to encounter a silent hardware fault at step  $k$ , i.e., instead of  $\mathbf{x}^{(k+1)}$  we could compute

$$\tilde{\mathbf{x}}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)}) + \tilde{\mathbf{x}}.$$

If  $\tilde{\mathbf{x}}^{(k+1)} \in \Gamma$ , Algorithm 1 will generate a new sequence  $\{\mathbf{y}^{(k)}\}_{k=0}^{\infty}$  where  $\mathbf{y}^{(0)} = \tilde{\mathbf{x}}^{(k+1)}$  and  $\mathbf{y}^{(i+1)} = \mathbf{G}(\mathbf{y}^{(i)})$ , and from Theorem 1  $\mathbf{y}^{(i)} \rightarrow \mathbf{x}_G$ . The total numerical error for the new sequence is bounded by

$$\|\mathbf{y}^{(i)} - \mathbf{x}_G\| \leq r^i \|\mathbf{y}^{(0)} - \mathbf{x}_G\| \leq r^i \|\mathbf{x}^{(k+1)} - \mathbf{x}_G\| + r^i \|\tilde{\mathbf{x}}\| \leq r^i r^{k+1} \|\mathbf{x}^{(0)} - \mathbf{x}_G\| + r^i \|\tilde{\mathbf{x}}\|.$$

The first term of the bound  $r^{i+k+1} \|\mathbf{x}^{(0)} - \mathbf{x}_G\|$  is identical to the one associated with the error free iteration. Therefore, the hardware error associated with the fixed point method that is caused by a single hardware failure is bounded by  $r^i \|\tilde{\mathbf{x}}\|$ , where  $i$  is the number of iterations after the failure and  $\|\tilde{\mathbf{x}}\|$  is the magnitude of the hardware induced perturbation. Analogously, if we encounter  $j$  hardware faults at iterations  $i_1, i_2, \dots, i_j$  with perturbations  $\tilde{\mathbf{x}}^{(i_1)}, \tilde{\mathbf{x}}^{(i_2)}, \dots, \tilde{\mathbf{x}}^{(i_j)}$ , then we have the following bound

$$\|\mathbf{x}^{(k)} - \mathbf{x}_G\| \leq r^k \|\mathbf{x}^{(0)} - \mathbf{x}_G\| + \sum_{l=1}^j r^{k-i_l} \|\tilde{\mathbf{x}}^{(i_l)}\|. \quad (3.3)$$

Thus, the hardware error associated with  $j$  silent faults is bounded by  $\sum_{l=1}^j r^{k-i_l} \|\tilde{\mathbf{x}}^{(i_l)}\|$ .

Next we consider the convergence of the method with respect to the hardware error as in Definition 2. The terms  $r^{k-i_l} \rightarrow 0$  as  $k \rightarrow \infty$ , however, for any non-zero  $r$ , the supremum in (2.1) will contain the unbounded terms  $\|\tilde{\mathbf{x}}^{(i_l)}\|$  and hence the expectation and variance of the hardware error will be unbounded. In practice, we expect to encounter perturbations of finite magnitude, however, a sufficiently large perturbation may require a prohibitive number of additional iterations to converge. Our earlier work shows how rescaling the entries of  $\mathbf{A}$  can be used to reduce the expected value of  $\|\tilde{\mathbf{x}}^{(i_l)}\|$  [10–12], however, not all problems can be efficiently rescaled and scaling cannot drive the expectation of  $\|\tilde{\mathbf{x}}^{(i_l)}\|$  to zero. Therefore, the fixed point iteration is not convergent with respect to silent faults if the evaluation of  $\mathbf{G}(\mathbf{x}^{(k)})$  is done in *fast* mode.

**3.2. Resilience enhancement techniques and convergence.** We assume that the evaluation of  $e^{(k)}$  and the conditional branching are computed in *safe* mode and we require that the expensive step that computes  $\mathbf{x}^{(k)}$  be executed in *fast* mode. The fixed point method has self correcting properties that can negate the effects of hardware error with small magnitude, however, a novel approach is needed to discriminate and reject iterates  $\mathbf{x}^{(k+1)}$  that are contaminated by hardware error with large magnitude.

First, note that the increments  $e^{(k)}$  defined in Algorithm 1 decay monotonically, i.e.,

$$e^{(k)} = \|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\| = \|\mathbf{G}(\mathbf{x}^{(k)}) - \mathbf{G}(\mathbf{x}^{(k-1)})\| \leq r\|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| = re^{(k-1)}. \quad (3.4)$$

The monotonic property (3.4) (i.e.,  $e_k \leq re_{k-1}$ ) can be easily verified at each step and deviation from this behavior is an indication of hardware malfunction. If the property fails, i.e., we observe something other than monotonic convergence, then we reject the next iterate  $\mathbf{x}^{(k+1)}$  and recompute  $\mathbf{G}(\mathbf{x}^{(k)})$ . In practice, the value of  $r$  may be unknown, hence we can introduce a tuning parameter  $\alpha$ , that is an upper estimate such that  $\alpha \in (r, 1]$ . If such an estimate is not available, we can start the iteration with  $\alpha = 1$  and from several increments  $e^{(k)}$  we can estimate  $\alpha$  as  $\alpha \approx e^{(k+1)}/e^{(k)}$ . Effectively, this will provide an upper bound on the hardware error introduced by the computations.

The monotonic condition (3.4) can be tested for any iterate  $k > 0$ , however, the first iterate needs a different choice for  $e^{(-1)}$ . An upper bound for  $e^{(0)}$  is given by

$$\begin{aligned} e^{(0)} = \|\mathbf{x}^{(1)} - \mathbf{x}^{(0)}\| &= \|\mathbf{x}^{(1)} - \mathbf{x}_G + \mathbf{x}_G - \mathbf{x}^{(0)}\| \\ &\leq r\|\mathbf{x}^{(0)} - \mathbf{x}_G\| + \|\mathbf{x}^{(0)} - \mathbf{x}_G\| \\ &= (r + 1)\|\mathbf{x}^{(0)} - \mathbf{x}_G\| \\ &\leq (\alpha + 1)\beta, \end{aligned}$$

where  $\beta$  is a second tuning parameter that is an upper bound for  $\|\mathbf{x}^{(0)} - \mathbf{x}_G\|$ . Thus, we set  $e^{(-1)} = (\alpha + 1)\beta$ , which is only needed for the first step of the iteration.

With all this in mind, we propose the following improved fixed point algorithm, where all steps except  $\mathbf{x}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)})$  are executed in safe mode:

**Algorithm 2** (Improved fixed point iteration).

*Given*  $\mathbf{G}$ , *initial*  $\mathbf{x}^{(0)}$ , *tolerance*  $\epsilon$  and  $k_{min}$   
*Select*  $\alpha \in (r, 1]$  and  $\beta > \|\mathbf{x}^{(0)} - \mathbf{x}_G\|$   
*Let*  $e^{(-1)} = (\alpha + 1)\beta$   
 $k = 0$   
**repeat**  
      $\mathbf{x}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)})$   
      $e^{(k)} = \|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|$   
     **if**  $e^{(k)} < \alpha e^{(k-1)}$  **then**  
         Accept  $\mathbf{x}^{(k+1)}$  (i.e.,  $k = k + 1$ )  
     **else**  
         Reject the step (i.e., make no modifications to  $k$ )  
     **end if**  
**until**  $(k < k_{min})$  **or**  $(e^{(k)} < \epsilon)$

The evaluation of  $\mathbf{G}(\mathbf{x}^{(k)})$  in Algorithm 2 is the only step that is susceptible to silent faults. We note that it is possible for a silent fault encountered at iteration  $k$  to result in failure of the monotonic condition  $e^{(k+2)} < \alpha e^{(k+1)}$  at step  $k+2$ , even if step  $k+1$  was computed without any faults. This delayed rejection can result in stagnation of Algorithm 2 and we address this issue in §3.4. First, we focus on an upper bound for the sum of numerical error and hardware errors.

We take an approach similar to [12]. Suppose we encounter a fault at iteration  $k$ , then we replace  $\mathbf{x}^{(k+1)}$  by

$$\tilde{\mathbf{x}}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)}) + \tilde{\mathbf{x}}.$$

For example, when  $k = 0$ , faulty initial iterate will be rejected and completely corrected unless

$$\|\mathbf{x}^{(1)} + \tilde{\mathbf{x}} - \mathbf{x}^{(0)}\| < (\alpha + 1)\beta.$$

On the other hand, if we accept a faulty iterate  $\tilde{\mathbf{x}}^{(1)}$ , we consider a new iteration  $\mathbf{y}^{(0)} = \tilde{\mathbf{x}}^{(1)}$  such that

$$\begin{aligned} \|\mathbf{y}_i - \mathbf{x}_G\| &\leq r^i \|\tilde{\mathbf{x}}^{(1)} - \mathbf{x}_G\| \\ &\leq r^i \|\tilde{\mathbf{x}}^{(1)} - \mathbf{x}^{(0)} + \mathbf{x}^{(0)} - \mathbf{x}_G\| \\ &\leq r^i \left( (\alpha + 1)\beta + \|\mathbf{x}^{(0)} - \mathbf{x}_G\| \right). \end{aligned} \quad (3.5)$$

It is important to note that the error bound for the new iteration, given by (3.5), is independent of the magnitude of the error  $\|\tilde{\mathbf{x}}\|$  introduced; and depends only on the parameters  $\alpha$ ,  $\beta$  and the accuracy of the initial iterate  $\|\mathbf{x}^{(0)} - \mathbf{x}_G\|$ . Furthermore, since  $r^i \rightarrow 0$ , the improved algorithm is convergent with respect to a fault in the first iteration (as  $i \rightarrow \infty$ ).

Next, we consider faults encountered at iterations  $k > 0$ . In order to account for the possibility of multiple faults, we employ an inductive argument that we base on the following lemma.

**Lemma 1** (Magnification of error). *Let  $\{\mathbf{x}^{(s)}\}_{s=0}^k$  be a sequence of iterates generated by Algorithm 2 and suppose that for some constant  $C > 0$  and some integer  $\tilde{k} \leq k$  the following holds*

$$\|\mathbf{x}^{(k)} - \mathbf{x}_G\| \leq Cr^{\tilde{k}} \quad \text{and} \quad \|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| \leq Cr^{\tilde{k}} \quad (3.6)$$

*If a fault is encountered during computation of  $\mathbf{x}^{(k+1)}$ , then for  $i > k$  (assuming no further faults are encountered) the following estimate hold*

$$\|\mathbf{x}^{(i)} - \mathbf{x}_G\| \leq (1 + \alpha)Cr^{i-k+\tilde{k}-1} \quad \text{and} \quad \|\mathbf{x}^{(i)} - \mathbf{x}^{(i-1)}\| \leq Cr^{i-k+\tilde{k}-1} \quad (3.7)$$

*Proof:* If we encounter a fault at iteration  $k > 0$ , Algorithm 2 will reject any perturbation  $\tilde{\mathbf{x}}^{(k)}$  unless

$$\|\mathbf{x}^{(k+1)} + \tilde{\mathbf{x}} - \mathbf{x}^{(k)}\| < \alpha \|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\|.$$

Rejection of the fault results in the fixed point iteration simply being delayed by one iteration, hence (3.7) holds trivially. If the fault is accepted, then we obtain a fault

$\tilde{\mathbf{x}}^{(k+1)} = \mathbf{x}^{(k+1)} + \mathbf{x}$  if the iteration proceeds without any further faults for  $i > k$

$$\begin{aligned}
 \|\mathbf{x}^{(i)} - \mathbf{x}_G\| &\leq r^{i-k-1} \|\tilde{\mathbf{x}}^{(k+1)} - \mathbf{x}_G\| \\
 &\leq r^{i-k-1} \|\tilde{\mathbf{x}}^{(k+1)} - \mathbf{x}^{(k)} + \mathbf{x}^{(k)} - \mathbf{x}_G\| \\
 &\leq r^{i-k-1} \alpha \|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| + r^{i-k-1} \|\mathbf{x}^{(k)} - \mathbf{x}_G\| \\
 &\leq C r^{i-k-1} \alpha r^{\tilde{k}} + C r^{i-k-1} r^{\tilde{k}} \\
 &= C(1 + \alpha) r^{i-k-1+\tilde{k}}
 \end{aligned}$$

Similarly,

$$\begin{aligned}
 \|\mathbf{x}^{(i)} - \mathbf{x}^{(i-1)}\| &\leq r^{i-k-1} \|\tilde{\mathbf{x}}^{(k+1)} - \mathbf{x}^{(k)}\| \\
 &\leq \alpha r^{i-k-1} \|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| \\
 &\leq \alpha C r^{i-k-1+\tilde{k}} \quad \square
 \end{aligned}$$

Combining (3.5) and Lemma (1) we now present the main result of this section

**Theorem 2** (Multiple faults). *Suppose that Algorithm 2 is executed for  $k$  iterations and it encounters  $j \leq k$  faults at iterations  $k_1 < k_2 < \dots < k_j$ . Then, at iteration  $k$*

$$\|\mathbf{x}^{(k)} - \mathbf{x}_G\| \leq C_{max} r^{k-j} (1 + \alpha)^j, \quad (3.8)$$

where

$$C_{max} = \max \left\{ \frac{(1+r)\|\mathbf{x}^{(0)} - \mathbf{x}_G\|}{r}, \frac{(\alpha+1)\beta + \|\mathbf{x}^{(0)} - \mathbf{x}_G\|}{r} \right\}$$

*Proof:* If  $k_1 > 0$  (i.e., the first iteration doesn't encounter faults), then for  $k \leq k_1$  assumption (3.6) holds with  $C = r^{(-1)}(1+r)\|\mathbf{x}^{(0)} - \mathbf{x}_G\|$  and  $\tilde{k} = k$ , and therefore, for  $k \geq k_1$  and  $k < k_2$  we have that

$$\|\mathbf{x}^{(k)} - \mathbf{x}_G\| \leq (1 + \alpha) C_{max} r^{k-1}, \quad \text{and} \quad \|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| \leq \alpha C_{max} r^{k-1}.$$

Conversely, if  $k_1 = 0$ , then according to (3.5), for  $0 < k \leq k_2$  assumption (3.6) holds with  $C = r^{-1}((\alpha+1)\beta + \|\mathbf{x}^{(0)} - \mathbf{x}_G\|)$  and  $\tilde{k} = k-1$ , and therefore, for  $k \geq k_2$  and  $k < k_3$

$$\|\mathbf{x}^{(k)} - \mathbf{x}_G\| \leq (1 + \alpha) C_{max} r^{k-2}, \quad \text{and} \quad \|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| \leq C_{max} r^{k-2}.$$

Inductively, suppose that for  $k \geq k_{j-1}$  and  $k < k_j$ , we have that assumption (3.6) holds with  $C = C_{max}(1 + \alpha)^{j-1}$  and  $\tilde{k} = k - j + 1$ , then according to Lemma 1 we have that for  $k \geq k_j$

$$\|\mathbf{x}^{(k)} - \mathbf{x}_G\| \leq (1 + \alpha)^j C_{max} r^{k-j}, \quad \text{and} \quad \|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| \leq \alpha(1 + \alpha)^{j-1} C_{max} r^{k-j} \quad \square$$

The result in Theorem 1 gives an estimate on the error associated with the resilient fixed point algorithm (i.e., Algorithm 2). More importantly, the estimate is independent of the magnitude of the introduced perturbations  $\tilde{\mathbf{x}}$  and the iterations  $(k_j)$  where the faults occur. Every time we encounter a fault, we waste one iteration (i.e., we decrement the exponent of  $r$  by one) and the fault introduces a magnification factor of  $1 + \alpha$ . This allows us to bound

the statistics of the error independent of  $\tilde{\mathbf{P}}$  and we need only focus our attention to the number of faults encountered in  $k$  iterations.

Each iteration of the fixed point scheme has approximately the same computational complexity, hence, we expect that the rate of faults, i.e., the Bernoulli parameters, associated with each step to be essentially constant. Let  $p_k = p$  for all iterations  $k$ , then the number of faults  $j$  encountered in  $k$  iterations follows a binomial distribution. Let  $\rho(j|k)$ , for  $0 \leq j \leq k$ , indicate the binomial probability measure, i.e., the probability of encountering  $j$  faults in  $k$  total iterations with fault rate  $p$ . Then, according to Theorem 2, for any distribution  $\tilde{\mathbf{P}}$  of the hardware induced perturbations, the statistical moments of the total numerical error can be bounded by:

$$\begin{aligned} \mathbb{E}_{\mathbf{p}, \tilde{\mathbf{P}}} [\|\mathbf{x}^{(k)} - \mathbf{x}_G\|] &\leq C_{max} \sum_{j=0}^k r^{k-j} (1 + \alpha)^j \rho(j|k), \\ \mathbb{V}_{\mathbf{p}, \tilde{\mathbf{P}}} [\|\mathbf{x}^{(k)} - \mathbf{x}_G\|] &\leq C_{max}^2 \sum_{j=0}^k r^{2k-2j} (1 + \alpha)^{2j} \rho(j|k). \end{aligned} \quad (3.9)$$

Therefore, the improved iteration is convergent if

$$\begin{aligned} \lim_{k \rightarrow \infty} \sum_{j=0}^k r^{k-j} (1 + \alpha)^j \rho(j|k) &= 0, \\ \lim_{k \rightarrow \infty} \sum_{j=0}^k r^{2k-2j} (1 + \alpha)^{2j} \rho(j|k) &= 0. \end{aligned} \quad (3.10)$$

Condition (3.10) is sufficient for convergence, so long as Algorithm 2 does not encounter a false positive rejection of a fault free iterate. If multiple faults are accepted in consecutive iterations, it is possible for Algorithm 2 to stagnate. We address this issue in §3.4, however, we first look at the rate of convergence of the modified algorithm.

**3.3. Convergence Rate.** The classical fixed point algorithm converges linearly, i.e.,  $O(r^k)$ , where  $r$  is the contraction constant (e.g., the spectral radius of the iteration matrix) and  $k$  is the number of iterations. However, in the presence of hardware faults, even if the algorithm produces an accurate approximation to the desired solution, the rate of convergence often times deteriorates [11, 15]. Furthermore, the slowdown is strongly influenced by the number of introduced hardware faults. Therefore, we are interested in the rate of convergence as a function of the rate of faults  $p$ .

The statistical moments of the hardware error are bounded in (3.9), where  $\rho(j|k)$  is the binomial measure

$$\rho(j|k) = \frac{k!}{j!(k-j)!} p^j (1-p)^{k-j}. \quad (3.11)$$

Then, by substituting (3.11) into (3.9), and combining like terms we obtain the estimates

$$\begin{aligned}
 \mathbb{E}_{\mathbf{p}, \tilde{\mathbf{P}}} \left[ \|\mathbf{x}^{(k)} - \mathbf{x}_G\| \right] &\leq C_{max} \sum_{j=0}^k \frac{k!}{j!(k-j)!} r^{k-j} (1+\alpha)^j p^j (1-p)^{k-j} \\
 &\leq C_{max} \sum_{j=0}^k \frac{k!}{j!(k-j)!} (r(1-p))^{k-j} ((1+\alpha)p)^j, \\
 \mathbb{V}_{\mathbf{p}, \tilde{\mathbf{P}}} \left[ \|\mathbf{x}^{(k)} - \mathbf{x}_G\| \right] &\leq C_{max}^2 \sum_{j=0}^k \frac{k!}{j!(k-j)!} r^{2k-2j} (1+\alpha)^{2j} p^j (1-p)^{k-j} \\
 &\leq C_{max}^2 \sum_{j=0}^k \frac{k!}{j!(k-j)!} (r^2(1-p))^{k-j} ((1+\alpha)^2 p)^j.
 \end{aligned}$$

Using the binomial theorem the expectation and variance are given by

$$\begin{aligned}
 \sup_{\tilde{\mathbf{P}}} \mathbb{E}_{\mathbf{p}, \tilde{\mathbf{P}}} \left[ \|\mathbf{x}^{(k)} - \mathbf{x}_G\| \right] &\leq C_{max} (r(1-p) + (1+\alpha)p)^k, \\
 \sup_{\tilde{\mathbf{P}}} \mathbb{V}_{\mathbf{p}, \tilde{\mathbf{P}}} \left[ \|\mathbf{x}^{(k)} - \mathbf{x}_G\| \right] &\leq C_{max}^2 (r^2(1-p) + (1+\alpha)^2 p)^k
 \end{aligned} \tag{3.12}$$

If  $(r(1-p) + (1+\alpha)p) < 1$  and  $(r^2(1-p) + (1+\alpha)^2 p) < 1$ , then the right hand side of (3.12) can be made arbitrarily small by increasing  $k$ , i.e., performing a sufficient number of iterations (extra work). Thus, Algorithm 2 will be convergent in terms of (2.1) if the following condition is satisfied

$$p < \frac{1-r}{(1+\alpha)-r} \quad \text{and} \quad p < \frac{1-r^2}{(1+\alpha)^2 - r^2}. \tag{3.13}$$

From (3.13) we get a relationship between the contraction properties of  $\mathbf{G}(\mathbf{x})$  and the reliability of the hardware, required to guarantee convergence of the resilient fixed point iteration. In particular, the convergence condition (3.13) allows us to quantify the following:

- (i) For perfectly reliable hardware, i.e.,  $p = 0$ , the exponential term of (3.12) becomes,

$$(r(1-p) + (1+\alpha)p)^k = r^k$$

Furthermore, the  $r^{-1}$  coefficient and the  $(\alpha+1)\beta + \|\mathbf{x}^{(0)} - \mathbf{x}_G\|$  terms in the definition of  $C_{max}$  manifest only at the occurrence of a fault, hence, when  $p = 0$  we can take  $C_{max} = \|\mathbf{x}^{(0)} - \mathbf{x}_G\|$  and thus recover the rate and constant associated with the convergence of the classical fixed point iteration.

- (ii) As the hardware fault rate  $p$  increases, then  $r(1-p) + (1+\alpha)p$  increases and the convergence deteriorates accordingly.
- (iii) As  $r \rightarrow 1$  and the conditioning of the fixed point method deteriorates, the convergence condition forces  $p \rightarrow 0$ , which implies problems with worse conditioning require more reliable hardware.

- (iv) The constant  $C_{max}$  is of order  $1/r$ , however, Theorem 1 states that error of magnitude  $1/r$  can be correct with a single (fault free) iteration. Thus, for a fixed  $p$ , as  $r \rightarrow 0$ , the  $C_{max}$  diverges but Algorithm 2 remain convergent.

Therefore, the modified fixed point iteration performs just as well as the classical method when executed on reliable hardware, while at the same time Algorithm 2 allows for convergence, even in the presence of silent hardware faults. Unreliable hardware would deteriorate the rate of convergence, which can be rigorously quantified by (3.12) and (3.13).

**3.4. Guarding Against False Positive Rejections.** The accept/reject criteria discussed above gives an upper bound for the statistics of the total numerical error. However, it is possible for a hardware fault to create a false-positive rejection of a following fault-free iteration. For example, suppose that at step  $k$  the result of  $\mathbf{G}(\mathbf{x}^{(k)})$  is perturbed so that

$$\tilde{\mathbf{x}} \approx -\mathbf{G}(\mathbf{x}^{(k)}), \quad \tilde{\mathbf{x}}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)}) + \tilde{\mathbf{x}} \approx \mathbf{x}^{(k)}.$$

Even though  $e^{(k)}$  is computed in safe mode, the perturbation  $\tilde{\mathbf{x}}$  will make  $e^{(k)}$  very small, which can cause either an early termination or any fault-free  $k+2$  iterate would no longer be accepted and Algorithm 2 will stagnate.

We propose a number of ways to address stagnation. The first possible approach is to apply redundancy techniques, where we assume that the probability of encountering two consecutive hardware failures that result in identical hardware error is negligible. This assumption is based on the consideration that every component of a supercomputer has a relatively low rate of failure and faults in different components are likely to result in very different hardware error. Suppose iteration  $k$  was rejected, then we store the computed presumably faulty  $\mathbf{x}^{(k+1)}$  and repeat the last iteration. If we reject the repeated iteration as well, then we compare the two rejected iterates, if they are within  $\epsilon$  of each other, then we overwrite the accept/reject criteria and force the acceptance of  $\mathbf{x}^{(k+1)}$  as the next iterate.

Another approach is to reject the last two iterates, upon encountering a fault, i.e., we reject both  $x_{k+1}$  and  $x_k$ . This is a very conservative approach that will guard against stagnation, however, every faulty iterate may result in the rejection of a fault free iterate as well. This is a feasible approach only if the rate of hardware faults is low and the increased storage cost of keeping two consecutive iterates is not of consideration.

Another redundancy approach would be to keep the history of the norms of the increments and test  $e_k$  against several of the past  $e_{k-1}, e_{k-2}, \dots$ . That is

```

if ( $e^{(k)} < \alpha e^{(k-1)}$ ) or ( $e^{(k)} < \alpha^2 e^{(k-2)}$ ) or ( $e^{(k)} < \alpha^3 e^{(k-3)}$ ) then
    Accept  $\mathbf{x}^{(k+1)}$ 
else
    Reject  $\mathbf{x}^{(k+1)}$ 
end if
    
```

While computationally and storage wise cheap, this approach will still fail if we encounter multiple consecutive faults causing stagnation, hence this method is only feasible if the probability of such event is negligible. Furthermore, if  $\alpha$  is an overestimate of  $r$ , then  $\alpha^3$  is an even bigger overestimate of  $r^3$  and hence this approach allows the acceptance of hardware error with larger magnitude.

In addition, there are multiple approaches to guard against early termination. Post-

processing in safe mode will guarantee that the residual associated with the final computed iterate is indeed within  $\epsilon$ . If the test fails, we can restart the fixed point iteration with the last iterate as the new initial guess  $\mathbf{x}^{(0)}$ . This is a very safe, but potentially expensive approach, as it requires one evaluation of (3.1) in safe mode.

Another approach is to require that two or more of the last computed increments  $e^{(k)}$  simultaneously satisfy the convergence criteria  $e^{(k)} < \epsilon$ , i.e., we replace the convergence test by

```
repeat
    ...
until ( $e^{(k)} < \epsilon$ ) and ( $e^{(k-1)} < \epsilon/r$ ) and ( $e^{(k-2)} < \epsilon/r^2$ ) ...
```

So long as encountering multiple consecutive faults that cause early termination has negligible probability, this approach will guarantee convergence.

All of above techniques are viable under different circumstances. We select the two that in the authors' opinion are most robust and most widely applicable. Specifically we consider the following Resilient fixed point Iteration.

**Algorithm 3** (Resilient fixed point Iteration).

*Given  $\mathbf{G}$ , initial  $\mathbf{x}_0$ , tolerance  $\epsilon$  and  $k_{min}$*   
*Let  $\alpha \in (r, 1]$ ,  $\beta > \|\mathbf{x}^{(0)} - \mathbf{x}_G\|$*   
*Let  $e^{(-1)} = (\alpha + 1)\beta$*   
 *$k = 0$*   
**repeat**  
      $\mathbf{x}^{(k+1)} = \mathbf{G}(\mathbf{x}^{(k)})$   
      $e^{(k)} = \|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|$   
     **if**  $e^{(k)} \leq \alpha e^{(k-1)}$  **then**  
         Accept  $\mathbf{x}^{(k+1)}$ , i.e.,  $k = k + 1$   
     **else**  
         **if** Last iterate was rejected and  $\|\mathbf{x}^{(k+1)} - \mathbf{x}_f\| \leq \epsilon$  **then**  
             Accept  $\mathbf{x}^{(k+1)}$  (i.e.,  $k = k + 1$ )  
         **else**  
             Reject the step, i.e., make no modifications to  $k$   
              $\mathbf{x}_f = \mathbf{x}^{(k+1)}$   
         **end if**  
     **end if**  
**until** ( $e^{(k)} < \epsilon$ ) and ( $e^{(k-1)} < \epsilon/\alpha$ )

**4. Numerical Example.** In this section, we verify the theoretical results by comparing the standard and resilient fixed point methods, when solving a parabolic partial differential equation (PDE). In particular, we consider the Jacobi method applied to the linear system of equations associated with one step of the implicit time integration of the heat equation. Let  $\Omega = [0, 1] \times [0, 1] \subset \mathbb{R}^2$  and consider the PDE

$$\begin{aligned} \frac{d}{d\tau} u(\tau, \xi, \eta) &= \frac{\partial^2}{\partial \xi^2} u(\tau, \xi, \eta) + \frac{\partial^2}{\partial \eta^2} u(\tau, \xi, \eta), & (\xi, \eta) \in \Omega, \tau > 0, \\ u(\tau, \xi, \eta)|_{\partial\Omega} &= 0, \\ u(0, \xi, \eta) &= \xi\eta(\xi - 1)(\eta - 1). \end{aligned}$$

We seek a numerical approximation to the solution  $u(\tau, \xi, \eta)$ . We discretize the problem in  $\Omega$  using a finite difference scheme with uniformly distributed nodes. Let

$$\{\xi_i\}_{i=1}^n, \quad \xi_i = \frac{i}{n+1}, \quad \{\eta_j\}_{j=1}^n, \quad \eta_j = \frac{j}{n+1},$$

and approximate

$$u(\tau, \xi_i, \eta_j) \approx u_{i,j}(\tau), \quad u_{i,j}(0) = \xi_i \eta_j (\xi_i - 1)(\eta_j - 1).$$

In this setting, the diffusion operator is discretized as

$$\begin{aligned} \frac{\partial^2}{\partial \xi^2} u_{i,j}(\tau) &\approx \frac{u_{i-1,j}(\tau) - 2u_{i,j}(\tau) + u_{i+1,j}(\tau)}{\Delta \xi^2}, \\ \frac{\partial^2}{\partial \eta^2} u_{i,j}(\tau) &\approx \frac{u_{i,j-1}(\tau) - 2u_{i,j}(\tau) + u_{i,j+1}(\tau)}{\Delta \eta^2}, \end{aligned}$$

where  $\Delta \xi = \Delta \eta = \frac{1}{n+1}$ , and  $u_{0,j}(\tau) = u_{n+1,j}(\tau) = u_{i,0}(\tau) = u_{i,n+1}(\tau) = 0$ . The spacial discretization results in  $n^2$  ordinary differential equations that can be written in a matrix form

$$\frac{d}{d\tau} \mathbf{v}(\tau) = \mathbf{L} \mathbf{v}(\tau),$$

where  $u_{i,j}(\tau)$  is given by  $v_m = u_{i,j}(\tau)$ , with  $m = (i-1)n + j$ , and  $\mathbf{L}$  is the matrix representation of the discretized Laplacian operator.

Finally, we evolve the system in time with the use of a backward Euler method. That is, we select a time step  $\Delta \tau$  and approximate

$$\mathbf{v}(\tau + \Delta \tau) \approx (\mathbf{I} - \Delta \tau \mathbf{L})^{-1} \mathbf{v}(\tau),$$

where  $\mathbf{I}$  is the identity matrix. At each time step we need to solve a system of linear equations. In order to study the resilience properties of the fixed point iteration, we consider only the first linear system (i.e.,  $\tau = 0$ ), such that

$$\mathbf{A} \mathbf{x} = \mathbf{b},$$

where  $\mathbf{A} = \mathbf{I} - \Delta \tau \mathbf{L}$ ,  $\mathbf{b} = \mathbf{v}(0)$  and  $\mathbf{x} \approx \mathbf{v}(\Delta \tau)$ .

**4.1. Simulating Hardware Faults.** We perform all computations on a desktop computer with Intel Sandy-Bridge-E 6-core CPU, and we implement our algorithms using MATLAB. The probability of a silent fault on the machine for the short duration of the computation is negligible. Nevertheless, we perform all tests twice to ensure consistency in the results. Since the hardware failure rate on this machine is low, we need to artificially introduce faults into the computations. At every iteration, we poll a pseudo random number using MATLAB's function `rand(1,1)`, which returns a uniformly distributed number in the range  $(0, 1)$ . If the random number is smaller than a specified threshold  $p$ , we introduce a perturbation to the evaluation of  $\mathbf{x}_{k+1}$ . We consider two distribution of the perturbations.

The first case, denoted by  $\tilde{\mathbf{P}}_u$ , uses the following formula

$$\tilde{\mathbf{x}} = 10^z \frac{1}{\|\mathbf{g}\|} \mathbf{g},$$

where  $\mathbf{g} \in \mathbb{R}^{n^2}$  is a vector with each component sampled from unit Gaussian distribution, i.e., MATLAB's function `randn(n*n,1)` and  $z \in [-9, 10]$  is uniformly distributed over the range, i.e., MATLAB's function `20*rand(1,1) - 9`. We have chosen this form of the perturbation because all directions in  $\mathbb{R}^{n^2}$  are well described, i.e.,  $\frac{1}{\|\mathbf{g}\|} \mathbf{g}$  is a unit vector with uniform distribution on the unit ball, and we have a representation of all the magnitudes in the range  $10^{-9}$  to  $10^{10}$ . While it is possible to encounter hardware fault that results in bigger error,  $10^{10}$  is sufficiently large to show the shortcomings of the classical fixed point method described in Algorithm 1, and larger error would be easily rejected by our resilient approach, i.e., Algorithm 3. On the other hand, hardware error that is significantly less than  $\epsilon$  will not have an effect on either algorithm. In what follows, we use  $\epsilon = 10^{-8}$  and hence  $10^{-9}$  is a suitable lower cutoff bound for the perturbations.

In the second case, we consider the worst possible scenario denoted by  $\tilde{\mathbf{P}}_w$ . In the case of linear fixed point methods (e.g., Jacobi method) the worst perturbation  $\tilde{\mathbf{x}}_w$  is aligned with the dominant characteristic direction of the iteration matrix and the magnitude  $\|\tilde{\mathbf{x}}_w\|$  is the largest magnitude that will be accepted by the test criteria, i.e.,  $\tilde{\mathbf{P}} = \delta_0(\tilde{\mathbf{x}} - \tilde{\mathbf{x}}_w)$ . Thus, the occurrence of faults remains random, however, the perturbations  $\tilde{\mathbf{x}}$  are the solutions to the deterministic optimization problems that maximizes  $\|\tilde{\mathbf{x}}\|$  in the direction of the dominant eigenvector of the iteration matrix, subject to the constraint that the perturbation has to be accepted by the test criteria of Algorithm 3.

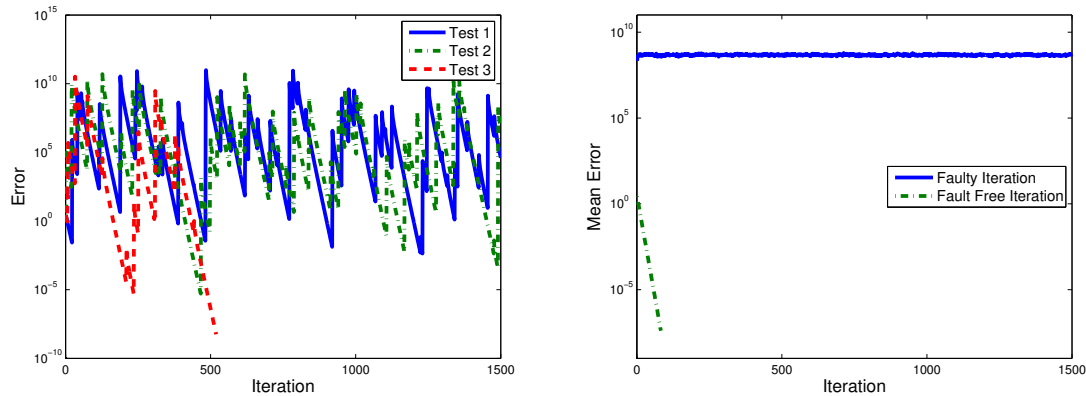
**4.2. Classical fixed point Iteration.** We first consider the classical fixed point method described in Algorithm 1. We take  $n = 100$ , which results in a system of 10,000 linear equations, i.e.,  $\mathbf{A} \in \mathbb{R}^{10,000 \times 10,000}$ . We consider the Jacobi fixed point method with  $\Delta t = 10^{-4}$ ,  $\mathbf{x}_0 = \mathbf{b}$  and  $\epsilon = 10^{-8}$  and the spectral radius of the resulting iteration matrix is  $r \approx 0.8028$ . We apply Algorithm 1 without any simulated hardware faults. In Figure 1 (dashed line on the right panel) we observe that the method converges linearly and it takes 83 iterations.

Next we want to observe the effects of the simulated hardware faults. We select the rate of silent faults to be one in ten, i.e.,  $p = 0.1$  and we use  $\tilde{\mathbf{P}}_u$ . This rate is extremely high compared to real world hardware, however, for the purpose of this study we want the expected number of faults to be significant in every realization.

The results from 4 simulations are shown in Figure 1 (left). In tests 1 and 2, the error does not drop below  $10^{-7}$  in the first 1500 iterations and the method fails to return an accurate answer. During test 3, the method reaches the tolerance in little over 500 iterations, after encountering a large number of consecutive fault-free iterations. Due to the large fault rate, encountering such a lucky run of no faults is a possible but improbable event and hence Algorithm 1 is unreliable.

In addition to the individual tests, we are also interested in the statistics of the total numerical error. We perform 10,000 tests and observe the average numerical error as a function of the number of iterations. Figure 1 (right) shows the lack of convergence for the first 1500 iterations. The expected value of the total numerical error depends on the rate

of faults and for a sufficiently small  $p$ , it is possible that the expectation can drop below  $\epsilon = 10^{-8}$  (e.g., see test case 3). Furthermore, if  $p$  is sufficiently small, the method will have a significant chance of not encountering any faults and returning an accurate approximation to the solution  $\mathbf{x}_G$ . However, the mean error will be bounded away from zero for any non-negative  $p$ . A resilient algorithm should provide results that are accurate in expectation for any arbitrary  $\epsilon$ , for as large of fault rate  $p$  as possible, therefore, according to Definition 2 the classical fixed point iteration does not converge with respect to hardware induced error.



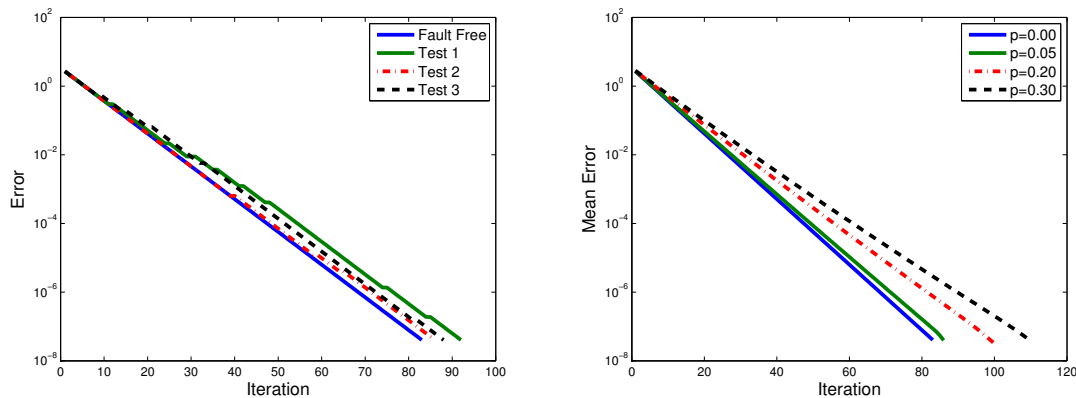
**Figure 1:** The Jacobi method implemented using the classical fixed point iteration of Algorithm 1 with simulated hardware faults at rate  $p = 0.1$ . **Left:** the result of four sample tests, where only Test 3 converged to desired tolerance but even that case took over 500 iterations. **Right:** (solid line) the expected value of the error as a function of the iteration computed by taking the average of 10,000 samples, i.e., with the use of Monte Carlo sampling. (dashed line) the method implemented without simulated faults converges in 82 iterations.

**4.3. Resilient fixed point Iteration.** We repeat the same set of experiments described in the previous section by keeping all of the parameters as described in section 4.2, except we replace the classic fixed point iteration with the resilient fixed point Algorithm 3. The tuning parameters that we use are  $\alpha = 1$  and  $\beta = 2\|\mathbf{b}\|$ , both of which are overestimate and hence conservatively safe. Figure 2 (left) shows the result of 3 tests compared to an error free iteration. In contrast to the classical method, the resilient iteration is much more reliable and despite the large number of simulated hardware faults, all tests converge in less than 100 iterations.

The test were performed with  $p = 0.2$ , which is in fact in violation of the convergence condition (3.13), which implies that the bound on  $p$  is conservative and hence (3.13) is sufficient but not necessary. This is due to the fact that the magnification factor  $(1 + \alpha)$  assumes that all hardware faults result in accepting a perturbation of maximum possible magnitude. In practice, many of the faulty iterates are either rejected or have a smaller magnitude and thus  $(1 + \alpha)$  is an overestimate. In Figure 2 (right), we plot the expected total numerical error as a function of the iteration for three different fault rates  $p = 0.05, 0.2, 0.3$ . While larger  $p$  implies that the algorithm would require more iterations to converge, the

resilient algorithm returns a reliable result even if 30% of the iterations are computed with faults. In Figure 3 (left), we plot just the mean for the  $p = 0.3$  case and the mean plus one standard deviation, demonstrating that we see convergence in both expectation and variance. Thus, according to Definition 2, the resilient fixed point method converges even in the presence of hardware faults.

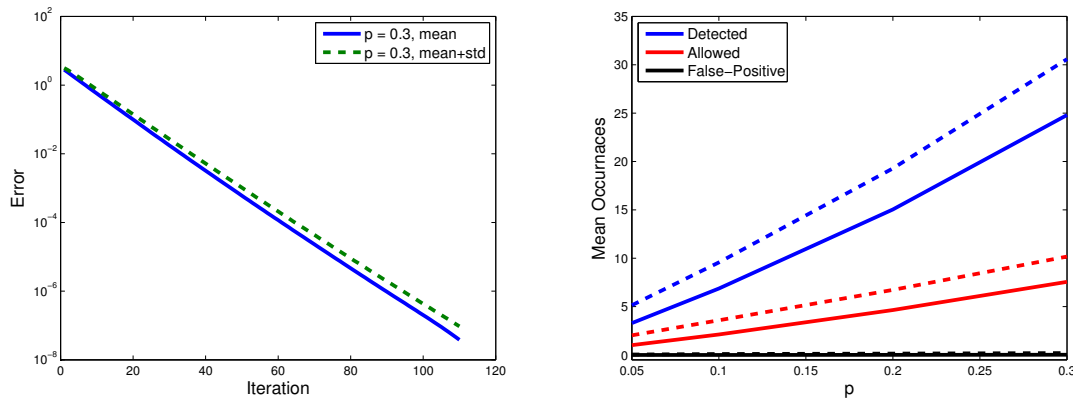
Additionally, Figure 3 (right) shows the mean number of faults that are detected (i.e., rejected), allowed (i.e., faults resulting in perturbations that is not rejected) and false-positive (i.e., rejections without a fault). The number of accepted and rejected faults increases with  $p$ , however, the number of false-positive remains very small for any  $p$ . Thus, the criteria defending against false-positive rejection rarely comes to use.



**Figure 2:** The Jacobi method implemented using the resilient fixed point iteration of Algorithm 3 with simulated hardware faults. **Left:** superimposed fault free iteration compared to three iterations using fault rate of  $p = 0.2$ . **Right:** average of the results using 10,000 simulations and observe convergence of the expected value of the total numerical error.

**4.4. Considering the Worst Case Scenario.** In this section we test the performance of the resilient Algorithm 3 using the worst case faults defined by  $\tilde{\mathbf{P}}_w$ . The purpose of the tests is to demonstrate the limitations of the method as well as the sharpness of the estimate (3.13).

We take  $N = 100$ ,  $\Delta\tau = 10^{-4}$  (resulting in  $r \approx 0.8028$ ),  $\alpha = 1$ ,  $\beta = 2\|\mathbf{b}\|$  and  $\epsilon = 10^{-8}$ . According to (3.13), the variance of the error associated with Algorithm 3 will converge for  $p < 0.1060$ , for  $p > 0.1060$  but  $p < 0.1647$  the mean error will tend to zero but the variance will not, for  $p > 0.1647$  the method will not be convergent. In Figure 4, we show the average results of 10,000 realization of Algorithm 3 with different values of  $p$ . Indeed, when  $p = 0.08$  we see convergence of mean and variance. For  $p = 0.11$ , we see convergence of the average error, but also large fluctuations of the mean indicating large variance, the problem with the variance becomes even more pronounced for  $p = 0.13$ . Note that on the semi-log plot we need large accuracy to obtain a *smooth* line, however, Monte Carlo accuracy is  $\sigma/\sqrt{S}$ , where  $\sigma$  is the standard deviation of the error for a given  $k$  and  $S$  is the total number of samples, thus, unless  $\sigma$  tends to zero as a function of  $k$ , we require a prohibitive number of samples to make a smooth plot. In the case of  $p = 0.14$ , the convergence is too slow for the



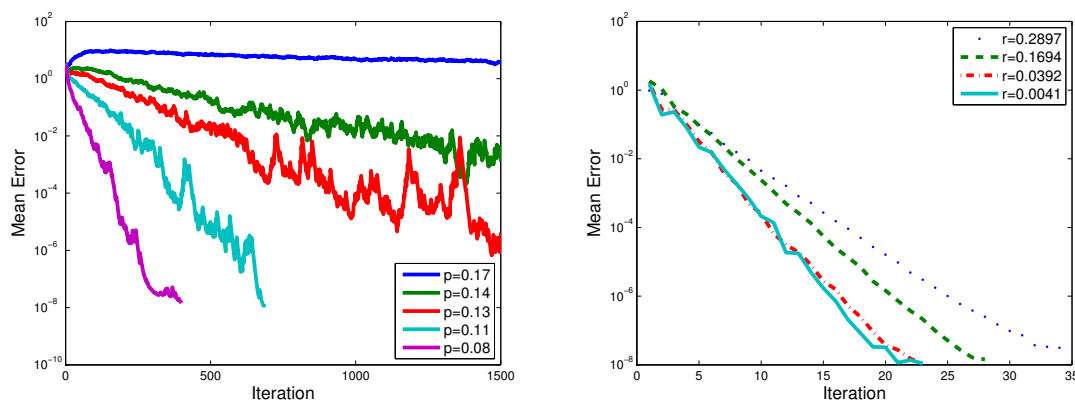
**Figure 3:** The Jacobi method implemented using the resilient fixed point iteration of Algorithm 3 with simulated hardware faults. **Left:** mean and mean plus one standard deviation of the error for the Jacobi method with  $p = 0.3$  **Right:** average number of detected, allowed, and false-positive faults observed in 10,000 simulations.

method to be practical. As soon as we violate condition (3.13) for the mean, at  $p = 0.17$ , Algorithm 3 is no longer convergent. This demonstrates the sharpness of our results.

Next, we fix  $N = 100$ ,  $p = 0.4$ ,  $\alpha = 0.4$ ,  $\beta = 2\|b\|$  and  $\epsilon = 10^{-8}$ , and we vary  $\Delta\tau$  making it smaller and smaller, which results in  $r \rightarrow 0$ . In the classical fixed point iteration, as  $r \rightarrow 0$ , the convergence rate increases and the method takes fewer and fewer iterations. However, according to (3.12), as  $r \rightarrow 0$  the term under the exponent will converge to  $(1+\alpha)p$  (or  $(1+\alpha)^2p$  for the variance), hence, at some point, we will not observe faster convergence even if we decrease  $r$ . Figure 4 (right) shows four examples corresponding to small  $r$ , and indeed we see very little difference between the convergence rate when  $r = 0.0392$  and  $r = 0.0041$  (corresponding to  $\Delta\tau = 5 \times 10^{-7}$  and  $\Delta\tau = 10^{-7}$ ). On the other hand, at  $k = 1$  there is little difference in the expected error and hence we do not observe the predicted  $1/r$  behavior of the constant  $C_{max}$ , which suggests that the estimate of the constant can be improved in future work.

**5. Conclusion.** This work demonstrates an analytic approach improving fixed point methods in the presence of silent hardware faults. Utilizing hardware and software based resilience techniques, we split our algorithm into *safe* and *fast* modes. While most of the computations are performed in the *fast* mode, a well chosen accept/reject criteria executed in *safe* mode can control the introduction and propagation of hardware induced error.

We extend the current framework of numerical analysis by removing the assumption that computations in fast mode can be performed reliably. We propose an architecture agnostic error model that can be used to analyze the hardware error propagation and an algorithm's convergence properties. In particular, we analyze the classical fixed point iteration and proposed several modifications that can dramatically improve resilience. We perform a numerical comparison between the standard and resilient fixed point methods. We demonstrate that the classical approach fails to converge in expectation, while the resilient method converges even when encountering a very high rate of silent hardware faults.



**Figure 4:** The Jacobi method implemented using the resilient fixed point iteration of Algorithm 3 with simulated hardware faults. **Left:** mean error for different values of  $p$  corresponding (bottom to top) to convergent behavior ( $p = 0.08$ ), convergent mean but not variance ( $p = 0.11, 0.13, 0.14$ ), and divergent behavior  $p = 0.17$ . **Right:** mean error for different values of  $r \rightarrow 0$ .

**Acknowledgements.** The authors were supported by the Office of Science of the US Department of Energy under grant number ERKJ245.

## REFERENCES

- [1] C. ANFINSON AND F. LUK, *Linear algebraic model of algorithm-based fault tolerance*, IEEE Transactions on Computers, 37 (1988), pp. 1599–1604.
- [2] A. AVIŽIENIS, *Fault-tolerance and fault-intolerance: Complementary approaches to reliable computing*, in ACM SIGPLAN Notices, vol. 10, ACM, 1975, pp. 458–464.
- [3] P. BRIDGES, M. HOEMMEN, K. FERREIRA, M. HEROUX, P. SOLTERO, AND R. BRIGHTWELL, *Cooperative application/OS DRAM fault recovery*, in Proceedings of the 2011 International Conference on Parallel Processing - Volume 2, Euro-Par’11, Berlin, Heidelberg, 2012, Springer-Verlag, pp. 241–250.
- [4] G. BRONEVETSKY AND B. DE SUPINSKI, *Soft error vulnerability of iterative linear algebra methods*, in Proceedings of the 22Nd Annual International Conference on Supercomputing, ICS ’08, New York, NY, USA, 2008, ACM, pp. 155–164.
- [5] F. CAPPELLO, A. GEIST, B. GROPP, L. KALE, B. KRAMER, AND M. SNIR, *Toward exascale resilience*, International Journal of High Performance Computing Applications, (2009).
- [6] M. CASAS, B. R. DE SUPINSKI, G. BRONEVETSKY, AND M. SCHULZ, *Fault resilience of the algebraic multi-grid solver*, in Proceedings of the 26th ACM International Conference on Supercomputing, ICS ’12, New York, NY, USA, 2012, ACM, pp. 91–100.
- [7] J. DONGARRA, J. HITTINGER, J. BELL, L. CHACON, R. FALGOUT, M. HEROUX, P. HOVLAND, E. NG, C. WEBSTER, AND S. WILD, *Applied mathematics research for exascale computing*. US Department of Energy Report, March 2014.
- [8] P. DU, A. BOUTELLER, G. BOSILCA, T. HERAULT, AND J. DONGARRA, *Algorithm-based fault tolerance for dense matrix factorizations*, in Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP ’12, New York, NY, USA, 2012, ACM, pp. 225–234.

- [9] J. ELLIOTT, M. HOEMMEN, AND F. MUELLER, *Evaluating the impact of SDC on the GMRES iterative solver*, in Proceedings of the 2014 IEEE 28th International Parallel and Distributed Processing Symposium, IPDPS '14, Washington, DC, USA, 2014, IEEE Computer Society, pp. 1193–1202.
- [10] ———, *Exploiting data representation for fault tolerance*, in Proceedings of the 5th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, Scala '14, Piscataway, NJ, USA, 2014, IEEE Press, pp. 9–16.
- [11] J. ELLIOTT, F. MUELLER, M. STOYANOV, AND C. WEBSTER, *Quantifying the impact of single bit flips on floating point arithmetic*, Tech. Rep. ORNL/TM-2013/282, Oak Ridge National Laboratory, One Bethel Valley Road, Oak Ridge, TN, 2013.
- [12] J. ELLIOTT, F. MUELLER, M. STOYANOV, AND C. WEBSTER, *Quantifying the impact of single bit flips on floating point arithmetic*, Tech. Rep. TR 2013-2, Dept. of Computer Science, North Carolina State University, Mar. 2013.
- [13] E. N. ELNOZAHY, L. ALVISI, Y.-M. WANG, AND D. B. JOHNSON, *A survey of rollback-recovery protocols in message-passing systems*, ACM Computing Surveys (CSUR), 34 (2002), pp. 375–408.
- [14] A. GEIST, *What is the monster in the closet?* Invited Talk at Workshop on Architectures I: Exascale and Beyond: Gaps in Research, Gaps in our Thinking, Aug. 2011.
- [15] M. HEROUX AND M. HOEMMEN, *Fault-tolerant iterative methods via selective reliability*, Tech. Rep. SAND2011-3915, Sandia National Laboratories, 2011.
- [16] K.-H. HUANG AND J. ABRAHAM, *Algorithm-based fault tolerance for matrix operations*, IEEE Transactions on Computers, C-33 (1984), pp. 518–528.
- [17] A. HWANG, I. STEFANOVICI, AND B. SCHROEDER, *Cosmic rays don't strike twice: Understanding the nature of DRAM errors and the implications for system design*, in Architectural Support for Programming Languages and Operating Systems, 2012, pp. 111–122.
- [18] J.-Y. JOU AND J. ABRAHAM, *Fault-tolerant matrix arithmetic and signal processing on highly concurrent computing structures*, Proceedings of the IEEE, 74 (1986), pp. 732–741.
- [19] A. ROY-CHOWDHURY, N. BELLAS, AND P. BANERJEE, *Algorithm-based error-detection schemes for iterative solution of partial differential equations*, Computers, IEEE Transactions on, 45 (1996), pp. 394–407.
- [20] W. RUDIN, *Principles of mathematical analysis*, vol. 3, McGraw-Hill New York, 1964.
- [21] P. SAO AND R. VUDUC, *Self-stabilizing iterative solvers*, in Proceedings of the Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems, Scala '13, New York, NY, USA, 2013, ACM, pp. 4:1–4:8.
- [22] B. SCHROEDER, E. PINHEIRO, AND W.-D. WEBER, *DRAM errors in the wild: A large-scale field study*, in SIGMETRICS Conference on Measurement and Modeling of Computer Systems, 2009, pp. 193–204.
- [23] J. SLOAN, R. KUMAR, AND G. BRONEVETSKY, *Algorithmic approaches to low overhead fault detection for sparse linear algebra*, in Dependable Systems and Networks (DSN), 2012 42nd Annual IEEE/IFIP International Conference on, IEEE, 2012, pp. 1–12.
- [24] M. SNIR, R. WISNIEWSKI, J. ABRAHAM, S. ADVE, S. BAGCHI, P. BALAJI, J. BELAK, P. BOSE, F. CAPPELLO, B. CARLSON, ET AL., *Addressing failures in exascale computing*, International Journal of High Performance Computing Applications, (2014), p. 1094342014522573.
- [25] V. SRIDHARAN AND D. LIBERTY, *A study of DRAM failures in the field*, in Supercomputing, Nov. 2012.
- [26] V. SRIDHARAN, J. STEARLEY, N. DEBARDELEBEN, S. BLANCHARD, AND S. GURUMURTHI, *Feng shui of supercomputer memory positional effects in DRAM and SRAM faults*, in High Performance Computing, Networking, Storage and Analysis (SC), 2013 International Conference for, IEEE, 2013, pp. 1–11.

