

Is the Multigrid Method Fault Tolerant? The Two-Grid Case



Approved for public release.
Distribution is unlimited.

Mark Ainsworth
Christian Glusa

June 2016

DOCUMENT AVAILABILITY

Reports produced after January 1, 1996, are generally available free via US Department of Energy (DOE) SciTech Connect.

Website <http://www.osti.gov/scitech/>

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
Telephone 703-605-6000 (1-800-553-6847)
TDD 703-487-4639
Fax 703-605-6900
E-mail info@ntis.gov
Website <http://www.ntis.gov/help/ordermethods.aspx>

Reports are available to DOE employees, DOE contractors, Energy Technology Data Exchange representatives, and International Nuclear Information System representatives from the following source:

Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831
Telephone 865-576-8401
Fax 865-576-5728
E-mail reports@osti.gov
Website <http://www.osti.gov/contact.html>

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, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents 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 or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Computer Science and Mathematics Division

**Is the Multigrid Method Fault Tolerant?
The Two-Grid Case**

Mark Ainsworth
ORNL Joint Faculty, Brown University

Christian Glusa
Brown University

Date Published:
June 2016

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831-6283
managed by
UT-BATTELLE, LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

IS THE MULTIGRID METHOD FAULT TOLERANT? THE TWO-GRID CASE*

MARK AINSWORTH [†] AND CHRISTIAN GLUSA [‡]

Abstract. The predicted reduced resiliency of next-generation high performance computers means that it will become necessary to take into account the effects of randomly occurring faults on numerical methods. Further, in the event of a hard fault occurring, a decision has to be made as to what remedial action should be taken in order to resume the execution of the algorithm. The action that is chosen can have a dramatic effect on the performance and characteristics of the scheme. Ideally, the resulting algorithm should be subjected to the same kind of mathematical analysis that was applied to the original, deterministic variant.

The purpose of this work is to provide an analysis of the behaviour of the multigrid algorithm in the presence of faults. Multigrid is arguably the method of choice for the solution of large-scale linear algebra problems arising from discretization of partial differential equations and it is of considerable importance to anticipate its behaviour on an exascale machine. The analysis of resilience of algorithms is in its infancy and the current work is perhaps the first to provide a mathematical model for faults and analyse the behaviour of a state-of-the-art algorithm under the model. It is shown that the Two Grid Method fails to be resilient to faults. Attention is then turned to identifying the minimal necessary remedial action required to restore the rate of convergence to that enjoyed by the ideal fault-free method.

Key words. Multigrid, Fault Tolerance, Resilience, Random Matrices, Convergence Analysis

AMS subject classifications. 65F10, 65N22, 65N55, 68M15

1. Introduction. President Obama’s executive order in the summer of 2015 establishing the National Strategic Computing Initiative¹ committed the US to the development of a *capable* exascale computing system. Given that the performance of the current number one machine Tianhe-2 is roughly one thirtieth of that of an exascale system, it is easy to underestimate the challenge posed by this task. One way to envisage the scale of the undertaking is that the combined processing power of the *entire* TOP 500 list is less than half of one exaflop (10^{18} floating point operations per second).

It is widely accepted that an exascale machine should respect a 20MW power envelope. Tianhe-2 already consumes 18MW of power, and if it were possible to simply upscale to exascale using the current technology, would require around 540MW or roughly the same amount of energy required to power half a million homes. In order to meet the power envelope, and other requirements, one is required to change the physical operating limits of the machine through such mechanisms as lower voltage logic thresholds, reduced cell capacitance and through minimisation of data movement. All of these factors contribute to a lower overall reliability of the machine in terms of random bit-flipping and corruption of logic states from extraneous sources such as cosmic rays and, failures of individual components. The issue is further exac-

*The efforts of Mark Ainsworth were partially supported by the SIRIUS (Science-driven Data Management for Multi-tiered Storage) project at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

[†]Division of Applied Mathematics, Brown University, 182 George St, Providence, RI 02912, USA. Joint Faculty with the Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA. (Mark_Ainsworth@Brown.edu)

[‡] Division of Applied Mathematics, Brown University, 182 George St, Providence, RI 02912, USA. (Christian.Glusa@Brown.edu)

¹Creating a National Strategic Computing Initiative, Executive Order No. 13702 of July 29th, 2015, Federal Register

erated by the additional numbers of hardware components required for an exascale system [8–10, 14].

The taxonomy of faults and failures by Avizienis et al. [3] distinguishes between *hard faults* and *soft faults*. Soft faults correspond to corruption of instructions as well as the data produced and used within the code but which allow the execution of the code to proceed albeit in a corrupted state. An example of such a fault would be flipping of individual bits in a floating point number. Some bit flips in the exponent, sign or in the significant components of the mantissa may result in a relatively large error that becomes detectable. Hard faults might correspond to a message from one or more compute nodes either being corrupted beyond all recognition (possibly resulting in an exception being thrown), or lost altogether in the event of a compute node failing completely. Hard faults result in the interruption of the execution of the actual code unless remedial action is taken. As such, hard faults constitute errors that are readily *detectable*. Although the likelihood q of a hard fault occurring can be assumed small, it is not negligible. The fact that an algorithm cannot continue after a hard fault without some kind of remedial action means that such faults cannot simply be ignored in the hope that the algorithm will recover naturally.

Most, if not all, existing algorithms were derived under the assumption that such faults cannot occur and accordingly, no indication is generally offered as to what might constitute an appropriate course of action should a fault be detected in a particular algorithm. Nevertheless, in the event of a hard fault occurring, a decision has to be made as to what remedial action should be taken in order to resume the execution of the algorithm. The action that is chosen can have a dramatic effect on the performance and characteristics of the algorithm and as such represents a vital part of the algorithm. Ideally, the resulting algorithm should be subjected to the same kind of mathematical analysis that was applied to the original, deterministic, variant of the algorithm. Historical considerations mean that such analysis is available for very few methods.

At the present time, the basic approach to dealing with faults consists of attempting to restore the algorithm to the state that would have existed had no fault occurred. Two common ways in which this might be accomplished include: *Checkpoint-restart* [20, 25] whereby the state of the system is stored (or checkpointed) at pre-determined intervals so that in the event of a fault occurring, the computation can be restarted from the stored state; or, *Process Replication* [20, 25], whereby each critical component of the overall process is replicated one or more times so that in the event of a component being subject to a fault, the true state can be restored by making reference to an unaffected replica. Hybrid variants of these strategies can also be contemplated. The drawbacks of such approaches are self-evident.

The purpose of this work is to provide an analysis of the behaviour of the multigrid algorithm on a machine prone to faults. The multigrid method is arguably the method of choice for the solution of large-scale linear algebra problems arising from discretization of partial differential equations. Consequently, it is of considerable importance to anticipate its behaviour on an exascale machine. The analysis of resilience of algorithms is in its infancy and the current work is perhaps the first to provide a mathematical model for faults and analyse the behaviour of a state-of-the-art algorithm under the model.

This work is organised as follows. In the next section we present a mathematical model for faults and a strategy for their mitigation. In Section 3 we will briefly introduce the multigrid algorithm and its variant when faults are taken into account. The random iteration matrix of the resulting *Fault-Prone Multigrid Method* is determined

and an appropriate metric for convergence is discussed. The main results of this work concerning the rate of convergence of the Fault-Prone Two Grid method is found in Section 4. Theorem 3 demonstrates that the Two Grid Method is not fault resilient. The minimal necessary remedial action is given by Theorem 4: protecting the prolongation restores the rate of convergence of the ideal fault-free method. Supporting numerical evidence is provided.

1.1. Related Work. Different techniques have been previously employed in order to achieve fault resilience for iterative methods in general and multigrid in particular. Replication was used in [11, 13], and checkpoint-restart in [7]. Stoyanov and Webster [26] proposed a method based on selective reliability for fixed point methods. Finally, Huber et al. [23] proposed a recovery method to mitigate the effect of hard faults.

2. Modelling Faults. In order to take account of the effect of detectable faults on the stability and convergence of a numerical algorithm, we develop a simple probabilistic model for this behaviour and describe how it is incorporated into the formulation of iterative algorithms. As a matter of fact, it will transpire that our model for detectable faults will also apply to the case of silent faults that are relatively small.

Our preferred approach to mitigating the effects of faults in multigrid and related iterative algorithms is rather simple: *the lost or corrupted result expected from a node is replaced by a zero*. That is to say, if a value $x \in \mathbb{R}$ is prone to possible hard faults, then we propose to continue the global computation using $\tilde{x} \in \mathbb{R}$ in place of x , where

$$\tilde{x} = \begin{cases} 0 & \text{if a fault is detected,} \\ x & \text{otherwise.} \end{cases}$$

This minimalist, *laissez-faire*, approach has obvious attractions in terms of restoring the system to a valid state without the need to (i) halt the execution (other than to perhaps migrate data to a substitute node in the event of a node failure); (ii) take or read any data checkpoints; (iii) recompute any lost quantities; or, (iv) compromise on resources through having to replicate processes. However, the efficacy of *laissez-faire* depends on the extent to which the convergence characteristics of the resulting algorithm mirror those of the fault-free deterministic variant. The subject of the present work is to carry out a detailed and rigorous mathematical analysis of this question in the context of the Two Grid Algorithm. The multigrid case will be considered in our subsequent work.

2.1. Probabilistic Model of Faults. In order to model the effects of the *laissez-faire* fault mitigation approach on an algorithm, we introduce a Bernoulli random variable given by

$$(1) \quad \chi = \begin{cases} 0 & \text{with probability } q, \\ 1 & \text{with probability } 1 - q. \end{cases}$$

If a scalar variable x is subject to faults, then the effect of the fault and the *laissez-faire* strategy is modelled by replacing the value of x by the new value $\tilde{x} = \chi x$. Evidently \tilde{x} is a random variable with mean and variance given by $\mathbb{E}[\tilde{x}] = (1 - q)x$ and $\mathbb{V}[\tilde{x}] = q(1 - q)x^2$. By the same token, the effect of a fault on a vector-valued variable $x \in \mathbb{R}^n$ is modelled in a similar fashion by defining $\tilde{x} = \mathcal{X}x$, where

$$\mathcal{X} = \text{diag}(\chi_1, \dots, \chi_n)$$

and χ_i are identically distributed Bernoulli random variables. The variables χ_i can be independent, thus modelling componentwise faults, or block-dependent, as would be the case for a node failure.

More formally, for given $\varepsilon > 0$, let \mathbb{S}_ε denote a set consisting of random matrices satisfying the following conditions:

ASSUMPTION 1.

1. Each $\mathbf{X} \in \mathbb{S}_\varepsilon$ is a random diagonal matrix.
2. For every $\mathbf{X} \in \mathbb{S}_\varepsilon$, there holds $\mathbb{E}[\mathbf{X}] = e(\mathbf{X})\mathbf{I}$, where $e(\mathbf{X}) > 0$, and $|e(\mathbf{X}) - 1| \leq C\varepsilon$ for some fixed $C > 0$.
3. For every $\mathbf{X} \in \mathbb{S}_\varepsilon$ there holds $\|\mathbb{V}[\mathbf{X}]\|_2 = \max_{i,j} |\text{Cov}(\mathbf{X}_{ii}, \mathbf{X}_{jj})| \leq \varepsilon$.

The parameter ε is regarded as being small meaning that the random matrices \mathbf{X} behave, with high probability, like an identity matrix. The matrices \mathbf{X} will appear at various points in our model for the Fault-Prone Multigrid Algorithm. It is easy to see that if \mathbf{X}_1 and \mathbf{X}_2 are independent diagonal matrices, then $\mathbf{X}_1\mathbf{X}_2$ is again a random diagonal matrix.

Sometimes, a random matrix will appear in two (or more) different places in a single equation and it will be important to clearly distinguish between the cases where (i) each of the two matrices is a *different* realisation of the random matrix, or (ii) the matrices both correspond to the *same* realisation of the random matrix. We shall adopt the convention whereby should a symbol appear twice, then (ii) holds: i.e. the two occurrences represent the *same realisation* of the random matrix and the matrices are therefore identical. However, if the square or higher power of a random matrix appears then (i) holds: i.e. each of the matrices in the product is a *different* realisation of the same random matrix.

2.2. Application to Silent Faults. As mentioned above, a typical example of a soft fault would be flipping of individual bits in a floating point number. Some bit flips in the exponent, sign or in the significant components of the mantissa may result in a relatively large error that becomes detectable. Such cases could be treated using the laissez-faire strategy described earlier. However, in other cases such faults may give a relatively small error and, as a result, be difficult (or impossible) to identify. Suppose that a vector x is subject to such a silent fault resulting in x being replaced by \tilde{x} in the machine. If the relative error is at most ε , then

$$(2) \quad \tilde{x}_i - x_i = \varepsilon_i \chi_i x_i$$

where χ_i is a Bernoulli random variable as in (1), and ε_i is a random variable on $[-\varepsilon, \varepsilon]$. Equally well, this means that

$$\tilde{x} = \mathbf{Y}x$$

where $\mathbf{Y} = \mathbf{I} + \text{diag}(\varepsilon_1\chi_1, \dots, \varepsilon_n\chi_n)$ is a random matrix that satisfies the previous conditions required for membership of \mathbb{S}_ε .

2.3. Scope of Faults Covered by the Analysis. Our analysis will cover faults that can be represented by a random matrix belonging to \mathbb{S}_ε . This means that the analysis will cover each of the cases (i) when the fault is detectable and mitigated using the laissez-faire strategy, and (ii) when the fault is silent but is relatively small in the sense that it may be modelled using (2). However, our analysis will not cover the important case involving faults which result in entries in the matrices or right hand sides in the problem being corrupted. Equally well, our analysis will not cover

the case of bit flips that result in a large relative error but which nevertheless remain undetected. Finally, this work is restricted to the solve phase of multigrid; we assume that the setup phase is protected from faults.

3. A Model for Fault-Prone Multigrid.

3.1. Multigrid Algorithm. Let \mathbf{A} be a symmetric, positive definite matrix arising from a finite element discretization of an elliptic partial differential equation in d spatial dimensions.

We wish to use a multigrid method to compute the solution of the problem

$$(3) \quad \mathbf{A}x = b$$

for a given load vector b . The multigrid method will utilise a nested hierarchy of progressively coarser grids of dimension $0 < n_0 < n_1 < \dots < n_L$. Restriction and prolongation operators are used to transfer vectors from one level in the hierarchy to the next:

$$\mathbf{R}_\ell^{\ell+1} : \mathbb{R}^{n_{\ell+1}} \rightarrow \mathbb{R}^{n_\ell}, \quad \mathbf{P}_{\ell+1}^\ell : \mathbb{R}^{n_\ell} \rightarrow \mathbb{R}^{n_{\ell+1}}.$$

It will be assumed that, as is often the case, these operators are related by the rule $\mathbf{R}_\ell^{\ell+1} = \left(\mathbf{P}_{\ell+1}^\ell\right)^T$. A sequence of matrices $\{\mathbf{A}_\ell\}_{\ell=0}^L$ on the coarser levels is defined recursively as follows

$$\mathbf{A}_L = \mathbf{A}; \quad \mathbf{A}_\ell = \mathbf{R}\mathbf{A}_{\ell+1}\mathbf{P}, \quad \ell = 0, \dots, L-1.$$

Here and in what follows, the indices on the prolongation and restriction are omitted whenever the appropriate choice of operator is clear from the context.

Smootheners are defined on each level $\ell = 1, \dots, L$ in the form

$$(4) \quad S_\ell(b_\ell, x_\ell) = x_\ell + \mathbf{N}_\ell(b_\ell - \mathbf{A}_\ell x_\ell),$$

where \mathbf{N}_ℓ is an approximation to the inverse of \mathbf{A}_ℓ . In particular, choosing the matrix $\mathbf{N}_\ell = \theta \mathbf{D}_\ell^{-1}$ corresponds to the damped Jacobi smoothener with damping factor θ .

The multigrid method is given in Algorithm 1 and may be invoked to obtain an approximate solution of problem (3) by a call to $MG_L(b, 0)$.

3.2. A Model for Fault-Prone Multigrid. The multigrid algorithm comprises a number of steps each of which may be affected by faults were the algorithm to be implemented on a fault-prone machine. In the absence of faults, a single iteration of multigrid replaces the current iterate x_ℓ by $MG_\ell(b_\ell, x_\ell)$ defined in Algorithm 1. However, in the presence of faults, a single iteration of multigrid means that x_ℓ is replaced by $\mathcal{M}G_\ell(b_\ell, x_\ell)$ where $\mathcal{M}G_\ell$ differs from MG_ℓ in general due to corruption of intermediate steps in the computations arising from faults. The object of this section is to develop a model for faults and define the corresponding operators $\mathcal{M}G_\ell$, $\ell \in \mathbb{N}$. For simplicity, we assume that the coarse grid is of moderate size meaning that the exact solve \mathbf{A}_0^{-1} on the coarsest grid is not prone to faults, i.e. $\mathcal{M}G_0(d_0, \cdot) = \mathbf{A}_0^{-1}d_0$. This is a reasonable assumption when the size of the coarse grid problem is sufficiently small that either faults are not an issue or, if they are, then replication can be used to achieve fault resilience. The extension of the analysis to the case of fault-prone coarse grid correction does not pose any fundamental difficulties but is not pursued in the present work.

Algorithm 1 Fault-free multigrid method MG_ℓ

Input: Right hand side b_ℓ ; Initial iterate x_ℓ


Output: $MG_\ell(b_\ell, x_\ell)$

```

1: if  $\ell = 0$  then return  $A_0^{-1}b_0$            // Exact solve on coarsest grid
2: else
3:   for  $i \leftarrow 1$  to  $\nu_1$  do
4:      $x_\ell \leftarrow S_\ell(b_\ell, x_\ell)$            //  $\nu_1$  pre-smoothing steps
5:    $d_{\ell-1} \leftarrow R(b_\ell - A_\ell x_\ell)$        // Restriction to coarser grid
6:    $e_{\ell-1} \leftarrow 0$ 
7:   for  $j \leftarrow 1$  to  $\gamma$  do
8:      $e_{\ell-1} \leftarrow MG_{\ell-1}(d_{\ell-1}, e_{\ell-1})$  //  $\gamma$  coarse grid correction steps
9:    $x_\ell \leftarrow x_\ell + P e_{\ell-1}$            // Prolongation to finer grid
10:  for  $i \leftarrow 1$  to  $\nu_2$  do
11:     $x_\ell \leftarrow S_\ell(b_\ell, x_\ell)$            //  $\nu_2$  post-smoothing steps

```

Global: $x_\ell \longrightarrow x_\ell + \mathcal{X}_\ell^{(S)} \mathbf{N}_\ell^{-1} (b_\ell - \mathbf{A}_\ell x_\ell)$

Distributed:  $[x_\ell]_1 \longrightarrow [x_\ell]_1 + \cancel{[\mathbf{N}_\ell^{-1}]_1 [b_\ell - \mathbf{A}_\ell x_\ell]_1} = [x_\ell]_1$

$[x_\ell]_2 \longrightarrow [x_\ell]_2 + [\mathbf{N}_\ell^{-1}]_2 [b_\ell - \mathbf{A}_\ell x_\ell]_2$

$[x_\ell]_3 \longrightarrow [x_\ell]_3 + [\mathbf{N}_\ell^{-1}]_3 [b_\ell - \mathbf{A}_\ell x_\ell]_3$

FIG. 1. Schematic representation of a node failure during smoothing and the remedial action taken by the laissez-faire approach in the case of three compute nodes. $[\bullet]_i$ represents the part of the quantity \bullet local to node i .

3.3. Smoothing. A single application of the smoothener S_ℓ to an iterate x_ℓ takes the form

$$(5) \quad S_\ell(b_\ell, x_\ell) = x_\ell + \mathbf{N}_\ell(b_\ell - \mathbf{A}_\ell x_\ell)$$

with each of the interior sub-steps in (5) being susceptible to faults.

The innermost step is the computation of the residual $\rho_\ell = b_\ell - \mathbf{A}_\ell x_\ell$. The action of the matrix \mathbf{A}_ℓ is applied repeatedly throughout the solution phase. We shall assume that neither the entries or structure of matrix \mathbf{A}_ℓ nor the right hand side b_ℓ are subject to corruption, only computation involving them. This would be the case were the information needed to compute the action of \mathbf{A}_ℓ placed in non-volatile random access memory (NVRAM) with relatively modest overhead. The net effect of faults in the computation of ρ_ℓ is that the preconditioner \mathbf{N}_ℓ does not act on the true residual but rather on a corrupted version modelled by $\mathcal{X}_1 \rho_\ell$ where \mathcal{X}_1 is a random diagonal matrix.

By the same token, the action of \mathbf{N}_ℓ is prone to faults, meaning that the true result $\mathbf{N}_\ell \mathcal{X}_1 \rho_\ell$ may be corrupted and is therefore modelled by $\mathcal{X}_2 \mathbf{N}_\ell \mathcal{X}_1 \rho_\ell$, where \mathcal{X}_2 is yet another random diagonal matrix. The matrix \mathbf{N}_ℓ corresponding to damped Jacobi is diagonal and hence $\mathcal{X}_2 \mathbf{N}_\ell \mathcal{X}_1 = \mathcal{X}^{(S)} \mathbf{N}_\ell$, where $\mathcal{X}^{(S)} = \mathcal{X}_1 \mathcal{X}_2$ is again

a random diagonal matrix. Consequently, the combined effect of the two sources of error can be modelled by a single random diagonal matrix.

In summary, our model for the action of a smoothener prone to faults consists of replacing the true smoothener S used in the pre- and post-smoothing step in the multigrid algorithm by the *non-deterministic smoothener*

$$(6) \quad S_\ell(b_\ell, x_\ell) = x_\ell + \mathcal{X}_\ell^{(S)} \mathbf{N}_\ell(b_\ell - \mathbf{A}_\ell x_\ell)$$

in which $\mathcal{X}_\ell^{(S)}$ is a random diagonal matrix which models the effect of the random faulty nature of the underlying hardware. The model (6) tacitly assumes that the current iterate x_ℓ remains fault-free. We illustrate the remedial action to a node failure in Figure 1.

3.4. Restriction, Prolongation and Coarse Grid Correction. The restriction of the residual described by the step

$$(7) \quad d_{\ell-1} = \mathbf{R}(b_\ell - \mathbf{A}_\ell x_\ell)$$

is prone to faults. Firstly, as in the case of the smoothener, the true residual ρ_ℓ is prone to corruption and is modelled by $\mathcal{X}_\ell^{(\rho)} \rho_\ell$. The resulting residual is then operated on by the restriction \mathbf{R} , which is itself prone to faults modelled using a random diagonal matrix $\mathcal{X}_{\ell-1}^{(R)}$. We arrive at the following model for the effect of faults on (7):

$$d_{\ell-1} = \mathcal{X}_{\ell-1}^{(R)} \mathbf{R} \mathcal{X}_\ell^{(\rho)} (b_\ell - \mathbf{A}_\ell x_\ell).$$

The coarse grid correction $e_{\ell-1}$ is obtained by performing γ iterations of $\mathcal{MG}_{\ell-1}$ with data $d_{\ell-1}$ and a zero initial iterate. The effect of faults when applying the prolongation to $e_{\ell-1}$ is also modelled by a random diagonal matrix $\mathcal{X}_\ell^{(P)}$ leading to the following model for the effect of faults on the coarse grid correction and prolongation steps:

$$x_\ell \leftarrow x_\ell + \mathcal{X}_\ell^{(P)} \mathbf{P} e_{\ell-1}.$$

3.5. Model for Multigrid Algorithm in Presence of Faults. Replacing each of the steps in the fault-free Multigrid Algorithm 1 with their non-deterministic equivalent yields the following model for the Fault-Prone Multigrid Algorithm 2 and defines the associated fault-prone multilevel operators \mathcal{MG}_ℓ , $\ell \in \mathbb{N}$.

The classical approach to the analysis of iterative solution methods for linear systems uses the notion of an iteration matrix. For example, in the case of the fault-free smoothening step (4)

$$(8) \quad x - S_\ell(\mathbf{A}_\ell x, y) = \mathbf{E}_\ell^S (x - y), \quad \forall x, y \in \mathbb{R}^{n_\ell}$$

where $\mathbf{E}_\ell^S = \mathbf{I} - \mathbf{N}_\ell \mathbf{A}_\ell$ is the iteration matrix. The corresponding result for the fault-prone smoothener (6) is given by

$$(9) \quad x - S_\ell(\mathbf{A}_\ell x, y) = \mathcal{E}_\ell^S (x - y), \quad \forall x, y \in \mathbb{R}^{n_\ell}$$

where the iteration matrix $\mathcal{E}_\ell^S = \mathbf{I} - \mathcal{X}_\ell^{(S)} \mathbf{N}_\ell \mathbf{A}_\ell$ is now *random*. By analogy with (8), we define the iteration matrix for the Fault-Prone Multigrid Algorithm 2 by the equation

$$(10) \quad x - \mathcal{MG}_\ell(\mathbf{A}_\ell x, y) = \mathcal{E}_\ell (x - y), \quad \forall x, y \in \mathbb{R}^{n_\ell}.$$

Algorithm 2 Model for Fault-Prone Multigrid Algorithm \mathcal{MG}_ℓ where $\mathcal{X}_\ell^{(\bullet)}$ are random diagonal matrices.

Input: Right hand side b_ℓ ; Initial iterate x_ℓ

Output: $\mathcal{MG}_\ell(b_\ell, x_\ell)$

```

1: if  $\ell = 0$  then return  $\mathbf{A}_0^{-1}b_0$            // Exact solve on coarsest grid
2: else
3:   for  $i \leftarrow 1$  to  $\nu_1$  do
4:      $x_\ell \leftarrow \mathcal{S}_\ell(b_\ell, x_\ell)$            //  $\nu_1$  pre-smoothing steps
5:      $d_{\ell-1} \leftarrow \mathcal{X}_{\ell-1}^{(R)} \mathbf{R} \mathcal{X}_\ell^{(\rho)} (b_\ell - \mathbf{A}_\ell x_\ell)$  // Restriction to coarser grid
6:      $e_{\ell-1} \leftarrow 0$ 
7:     for  $j \leftarrow 1$  to  $\gamma$  do
8:        $e_{\ell-1} \leftarrow \mathcal{MG}_{\ell-1}(d_{\ell-1}, e_{\ell-1})$  //  $\gamma$  coarse grid correction steps
9:        $x_\ell \leftarrow x_\ell + \mathcal{X}_\ell^{(P)} \mathbf{P} e_{\ell-1}$  // Prolongation to finer grid
10:    for  $i \leftarrow 1$  to  $\nu_2$  do
11:       $x_\ell \leftarrow \mathcal{S}_\ell(b_\ell, x_\ell)$            //  $\nu_2$  post-smoothing steps

```

In particular, $x - \mathcal{MG}_0(\mathbf{A}_0 x, \bullet) = 0$ and hence $\mathcal{E}_0 = \mathbf{0}$, i.e. the zero matrix.

The ν_1 pre-smoothing steps in Algorithm 2 can be expressed in the form

$$x^{(0)} = x_\ell; \quad x^{(i)} = \mathcal{S}_\ell(b_\ell, x^{(i-1)}), \quad i = 1, \dots, \nu_1.$$

Using (9) gives

$$x - x_\ell^{(i)} = \mathcal{E}_\ell^S(x - x_\ell^{(i-1)}), \quad i = 1, \dots, \nu_1$$

and hence

$$(11) \quad x - x_\ell^{(\nu_1)} = \left(\mathcal{E}_\ell^S\right)^{\nu_1} (x - x_\ell).$$

Using this notation means that the vector $d_{\ell-1}$ appearing in Algorithm 2 is given by

$$d_{\ell-1} = \mathcal{X}_{\ell-1}^{(R)} \mathbf{R} \mathcal{X}_\ell^{(\rho)} (b_\ell - \mathbf{A}_\ell x_\ell^{(\nu_1)}) = \mathcal{X}_{\ell-1}^{(R)} \mathbf{R} \mathcal{X}_\ell^{(\rho)} \mathbf{A}_\ell (x - x_\ell^{(\nu_1)}).$$

where $b_\ell = \mathbf{A}_\ell x$.

The coarse grid correction steps of Algorithm 2 can be written in the form

$$e_{\ell-1}^{(0)} = 0; \quad e_{\ell-1}^{(j)} = \mathcal{MG}_{\ell-1}(d_{\ell-1}, e_{\ell-1}^{(j-1)}), \quad j = 1, \dots, \gamma.$$

Let $z = \mathbf{A}_{\ell-1}^{-1} d_{\ell-1}$ then

$$z - e_{\ell-1}^{(j)} = z - \mathcal{MG}_{\ell-1}(d_{\ell-1}, e_{\ell-1}^{(j-1)}) = \mathcal{E}_{\ell-1}(z - e_{\ell-1}^{(j-1)}), \quad j = 1, \dots, \gamma$$

thanks to (10). Iterating this result and recalling that $e_{\ell-1}^{(0)} = 0$, gives $z - e_{\ell-1}^{(\gamma)} = \mathcal{E}_{\ell-1}^\gamma z$, or, equally well,

$$e_{\ell-1}^{(\gamma)} = (\mathbf{I} - \mathcal{E}_{\ell-1}^\gamma) \mathbf{A}_{\ell-1}^{-1} d_{\ell-1}.$$

Using these notations, the prolongation to the finer grid step in Algorithm 2 takes the form

$$x_\ell^{(P)} = x_\ell^{(\nu_1)} + \mathcal{X}_\ell^{(P)} \mathbf{P} e_{\ell-1}^{(\gamma)}$$

and hence, by collecting results, we deduce that

$$\begin{aligned} x - x_\ell^{(P)} &= (x - x_\ell^{(\nu_1)}) - \mathcal{X}_\ell^{(P)} \mathbf{P} e_{\ell-1}^{(\gamma)} \\ &= \left[\mathbf{I} - \mathcal{X}_\ell^{(P)} \mathbf{P} (\mathbf{I} - \mathcal{E}_{\ell-1}^\gamma) \mathbf{A}_{\ell-1}^{-1} \mathcal{X}_{\ell-1}^{(R)} \mathbf{R} \mathcal{X}_\ell^{(\rho)} \mathbf{A}_\ell \right] (x - x_\ell^{(\nu_1)}). \end{aligned}$$

Arguments identical to those leading to (11) give the following result

$$x - \mathcal{MG}_\ell(\mathbf{A}_\ell x, x_\ell) = \left(\mathcal{E}_\ell^S \right)^{\nu_2} (x - x_\ell^{(P)}).$$

which, in view of identity (10) and (11), yields the following recursive formula for the iteration matrix of the Fault-Prone Multigrid Algorithm 2:

$$(12) \quad \mathcal{E}_\ell = \left(\mathcal{E}_\ell^{S, \text{post}} \right)^{\nu_2} \left[\mathbf{I} - \mathcal{X}_\ell^{(P)} \mathbf{P} (\mathbf{I} - \mathcal{E}_{\ell-1}^\gamma) \mathbf{A}_{\ell-1}^{-1} \mathcal{X}_{\ell-1}^{(R)} \mathbf{R} \mathcal{X}_\ell^{(\rho)} \mathbf{A}_\ell \right] \left(\mathcal{E}_\ell^{S, \text{pre}} \right)^{\nu_1},$$

for $\ell = 1, \dots, L$ with $\mathcal{E}_0 = \mathbf{0}$. Here, we have used superscripts pre and post to reflect that the pre- and post-smootheners are independent realisations of the same random matrix.

By setting $\mathcal{E}_{L-1} = \mathbf{0}$ and applying the recursion (12) in the case $\ell = L$ yields a formula for the iteration matrix of the *Fault-Prone Two Grid Algorithm*:

$$\begin{aligned} (13) \quad \mathcal{E}_L^{TG} &= \left(\mathcal{E}_L^{S, \text{post}} \right)^{\nu_2} \left[\mathbf{I} - \mathcal{X}_L^{(P)} \mathbf{P} \mathbf{A}_{L-1}^{-1} \mathcal{X}_{L-1}^{(R)} \mathbf{R} \mathcal{X}_L^{(\rho)} \mathbf{A}_L \right] \left(\mathcal{E}_L^{S, \text{pre}} \right)^{\nu_1} \\ &= \left(\mathcal{E}_L^{S, \text{post}} \right)^{\nu_2} \mathcal{E}_L^{CG} \left(\mathcal{E}_L^{S, \text{pre}} \right)^{\nu_1}, \end{aligned}$$

corresponding to using an exact solver on level $L-1$. Here \mathcal{E}_L^{CG} is the iteration matrix of the exact fault-prone coarse grid correction.

3.6. Lyapunov Spectral Radius and Replica Trick. If the Fault-Prone Multigrid, Algorithm 2, is applied using a starting iterate with error $e^{(0)}$, then the error after $N \in \mathbb{N}$ iterations is given by $e^{(N)} = \mathcal{E}_L^N e^{(0)}$. The matrix \mathcal{E}_L defined recursively by (12) is random, and the product \mathcal{E}_L^N should be interpreted as a product of N independent samples of the matrix \mathcal{E}_L . The Fault-Prone Multigrid iteration will converge provided that

$$\lim_{N \rightarrow \infty} \left\| \mathcal{E}_L^N e^{(0)} \right\|^{1/N} < 1 \quad \text{a.s.}$$

where $\|\bullet\|$ is any norm on the (finite dimensional) space \mathbb{R}^{n_L} . The limit is understood in the sense of almost sure *a.s.* convergence for a given initial iterate $e^{(0)}$. A classic result, due to Furstenberg and Kesten [16], states that if $\mathbb{E} \left[\log (\mathcal{E}_L)^+ \right] < \infty$ then the following limit exists

$$\lim_{N \rightarrow \infty} \left\| \mathcal{E}_L^N \right\|^{1/N} = \varrho(\mathcal{E}_L) \quad \text{a.s.}$$

where the non-random quantity

$$\varrho(\mathcal{E}_L) = \lim_{N \rightarrow \infty} \exp \left\{ \mathbb{E} \left[\log \left\| \mathcal{E}_L^N \right\| \right]^{1/N} \right\}$$

is the *Lyapunov spectral radius*. Further details and properties of the Lyapunov spectral radius will be found in the books by Bougerol and Lacroix [4] or Crisanti et al. [12].

Suppose the matrix \mathcal{E}_L is non-random, then the Lyapunov spectral radius coincides with the usual spectral radius $\rho(\mathcal{E}_L)$, and it extends the notion of an asymptotic rate of convergence to random iterative schemes. In particular, $\varrho(\mathcal{E}_L) < 1$ means that the matrix product is almost surely convergent, whereas $\varrho(\mathcal{E}_L) > 1$ means that it is almost surely divergent.

Consequently, the convergence analysis of the Fault-Prone Multigrid Method, Algorithm 2, boils down to the study of the Lyapunov spectral radius of the random iteration matrix \mathcal{E}_L . Unfortunately, the treatment of the Lyapunov spectral radius is considerably more complex than dealing with the usual spectral radius. In general its determination is a NP-hard problem [28]. Since the Lyapunov spectral radius is not necessarily a continuous function of the fault probabilities [1], straightforward perturbation analysis cannot be employed. The following result is sometimes referred to as the *Replica Trick*:

THEOREM 1 (Replica trick [12]). *Let \mathcal{E} be a random square matrix. Then*

$$\varrho(\mathcal{E}) \leq \sqrt{\rho(\mathbb{E}[\mathcal{E} \otimes \mathcal{E}])},$$

where $\rho(\bullet)$ denotes the usual spectral radius.

We refer the reader to Appendix A for the proof of Theorem 1. One attractive feature of this estimate is the appearance of the usual spectral radius of the non-random matrix $\mathbb{E}[\mathcal{E}_L \otimes \mathcal{E}_L]$ rather than the Lyapunov spectral radius which, however, comes at the price of having to deal with the Kronecker product $\mathcal{E}_L \otimes \mathcal{E}_L$ which, as we shall see later, presents its own difficulties.

3.7. Behaviour of the Smoother under the Model. The iteration matrix for the fault-prone smoother (6) is given by

$$\mathcal{E}_\ell^S = \mathbf{I} - \mathcal{X}_\ell^{(S)} \mathbf{N}_\ell \mathbf{A}_\ell,$$

Here, we take $\mathcal{X}_\ell^{(S)} = \text{diag}(\chi_1, \dots, \chi_{n_\ell})$ to be a random diagonal matrix of componentwise faults with $\mathbb{E}(\chi_j) = 1 - q$ and $\text{Cov}(\chi_j, \chi_j) = q(1 - q)\delta_{ij}$. The behaviour of the fault-prone smoother is governed by the Lyapunov spectral radius $\varrho(\mathcal{E}_\ell^S)$.

THEOREM 2. *Let $\mathbf{E}_\ell^S = \mathbf{I} - \mathbf{N}_\ell \mathbf{A}_\ell$ be the iteration matrix for any convergent smoother. Then, the corresponding fault-prone smoother satisfies*

$$\varrho(\mathcal{E}_\ell^S) \leq \sqrt{(1 - q) \left\| \mathbf{E}_\ell^S \right\|_2^2 + q \left(\left\| \mathbf{E}_\ell^S \right\|_2 + \left\| \mathbf{N}_\ell \mathbf{A}_\ell \right\|_2 \right)^2},$$

for all $q \in [0, 1]$.

Proof. Let $\mathcal{B} = \mathcal{X}^{(S)} \mathbf{N} \mathbf{A}$ and dispense with subscripts and superscripts for the duration of the proof. Then

$$\begin{aligned} \mathbb{E}[\mathcal{E} \otimes \mathcal{E}] &= \mathbb{E}[\mathbf{I} \otimes \mathbf{I} - \mathcal{B} \otimes \mathbf{I} - \mathbf{I} \otimes \mathcal{B} + \mathcal{B} \otimes \mathcal{B}] \\ &= \mathbf{I} \otimes \mathbf{I} - \mathbb{E}[\mathcal{B}] \otimes \mathbf{I} - \mathbf{I} \otimes \mathbb{E}[\mathcal{B}] + \mathbb{E}[\mathcal{B}] \otimes \mathbb{E}[\mathcal{B}] \\ &\quad - \mathbb{E}[\mathcal{B}] \otimes \mathbb{E}[\mathcal{B}] + \mathbb{E}[\mathcal{B} \otimes \mathcal{B}] \\ (14) \quad &= \mathbb{E}[\mathcal{E}] \otimes \mathbb{E}[\mathcal{E}] + \mathbb{V}[\mathcal{B}] \end{aligned}$$

and

$$\mathbb{V}[\mathcal{B}] = \mathbb{V}[\mathcal{X}^{(S)}] (\mathbf{N} \mathbf{A}) \otimes (\mathbf{N} \mathbf{A}) = q(1 - q) \mathbf{K}(\mathbf{N} \mathbf{A}) \otimes (\mathbf{N} \mathbf{A}),$$

where

$$\mathbf{K} = \text{blockdiag} \left(\bar{\mathbf{e}}^{(j)} \otimes \left(\bar{\mathbf{e}}^{(j)} \right)^T, j = 1, \dots, n \right)$$

and $\bar{\mathbf{e}}^{(j)}$ the j -th canonical unit vector in \mathbb{R}^n . Hence, since $\|\mathbf{C} \otimes \mathbf{C}\|_2 = \|\mathbf{C}\|_2^2$ and $\|\mathbf{K}_\ell \mathbf{C}\|_2 \leq \|\mathbf{C}\|_2$ for any compatible matrix \mathbf{C} , we obtain

$$\|\mathbb{V}[\mathbf{B}]\|_2 = q(1-q) \|\mathbf{K}(\mathbf{N}\mathbf{A}) \otimes (\mathbf{N}\mathbf{A})\|_2 \leq q(1-q) \|\mathbf{N}\mathbf{A}\|_2^2$$

and

$$\|\mathbb{E}[\mathbf{E}] \otimes \mathbb{E}[\mathbf{E}]\|_2 = \|\mathbb{E}[\mathbf{E}]\|_2^2 = \|\mathbf{I} - (1-q)\mathbf{N}\mathbf{A}\|_2^2.$$

Consequently,

$$\begin{aligned} \|\mathbb{E}[\mathbf{E} \otimes \mathbf{E}]\|_2 &\leq \|\mathbf{I} - (1-q)\mathbf{N}\mathbf{A}\|_2^2 + q(1-q) \|\mathbf{N}\mathbf{A}\|_2^2 \\ &= \|\mathbf{E} + q\mathbf{N}\mathbf{A}\|_2^2 + q(1-q) \|\mathbf{N}\mathbf{A}\|_2^2 \\ &\leq (\|\mathbf{E}\|_2 + q\|\mathbf{N}\mathbf{A}\|_2)^2 + q(1-q) \|\mathbf{N}\mathbf{A}\|_2^2 \\ &= (1-q) \|\mathbf{E}\|_2^2 + q(\|\mathbf{E}\|_2 + \|\mathbf{N}\mathbf{A}\|_2)^2 \end{aligned}$$

and the result follows thanks to Theorem 1. \square

In order to illustrate Theorem 2, we consider the situation where the smoothener is taken to be a convergent relaxation scheme so that $\|\mathbf{E}_\ell^S\|_2 = \gamma \in [0, 1)$. The triangle inequality gives $\|\mathbf{N}_\ell \mathbf{A}_\ell\|_2 \leq 1 + \gamma$ and hence, thanks to Theorem 2, we deduce that

$$\varrho(\mathbf{E}_\ell^S) \leq \sqrt{\gamma^2 + q(1 + 4\gamma + 3\gamma^2)}.$$

As a consequence, we find that a sufficient condition for the fault-prone smoothener to remain convergent is that the error rate is sufficiently small: $q < (1 - \gamma)/(1 + 3\gamma)$.

4. Summary of Main Results. Let the following standard assumptions for the convergence of the fault-free multigrid method be satisfied [5, 18, 19, 27]:

ASSUMPTION 2 (Smoothing property). *There exists $\eta : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ satisfying $\lim_{\nu \rightarrow \infty} \eta(\nu) = 0$ and such that*

$$\|\mathbf{A}_\ell (\mathbf{E}_\ell^S)^\nu\|_2 \leq \eta(\nu) \|\mathbf{A}_\ell\|_2, \quad \nu \geq 0, \ell = 1, \dots, L.$$

ASSUMPTION 3 (Approximation property). *There exists a constant C_A such that*

$$\|\mathbf{E}_\ell^{CG} \mathbf{A}_\ell^{-1}\|_2 \leq \frac{C_A}{\|\mathbf{A}_\ell\|_2}, \quad \ell = 1, \dots, L.$$

ASSUMPTION 4. *The smoothener is non-expansive, i.e. $\rho(\mathbf{E}_\ell^S) = \|\mathbf{E}_\ell^S\|_A \leq 1$, and there exists a non-increasing function $C_S : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ such that*

$$\|(\mathbf{E}_\ell^S)^\nu\|_2 \leq C_S(\nu) \quad \nu \geq 1, \ell = 1, \dots, L.$$

ASSUMPTION 5. *There exist positive constants \underline{C}_p and \overline{C}_p such that*

$$\underline{C}_p^{-1} \|x_\ell\|_2 \leq \|\mathbf{P}x_\ell\|_2 \leq \overline{C}_p \|x_\ell\|_2 \quad \forall x_\ell \in \mathbb{R}^{n_\ell}, \ell = 0, \dots, L-1.$$

Assumptions 2 and 3 guarantee convergence of the Two Grid Method, i.e. $\|\mathbf{E}_\ell^{TG}\|_2 \leq C < 1$, whereas Assumptions 4 and 5 can be used to show multigrid convergence with respect to $\|\bullet\|_2$.

4.1. Analysis of the Fault-Prone Two Grid Method. We first analyse the Fault-Prone Two Grid Method in the following setting: Let $\Omega \subset \mathbb{R}^d$ be a convex polyhedral domain, $\Gamma_D \subseteq \partial\Omega$ be non-empty and $V = \{v \in H^1(\Omega) : v = 0 \text{ on } \Gamma_D\}$. Let \mathcal{T}_0 be a partitioning of Ω into the union of elements each of which is a $d + 1$ -simplex, such that the intersection of any pair of distinct elements is a single common vertex, edge, face or sub-simplex of both elements. For $\ell \in \mathbb{N}$, the partition \mathcal{T}_ℓ is obtained by uniformly subdividing each element in $\mathcal{T}_{\ell-1}$ into subsimplices such that \mathcal{T}_ℓ satisfies the same conditions as \mathcal{T}_0 [24]. Let \mathbb{P}_k , $k \in \mathbb{N}$, consist of polynomials on \mathbb{R}^d of total degree at most k . The subspace V_ℓ consists of continuous piecewise polynomials belonging to \mathbb{P}_k on each element on the partition \mathcal{T}_ℓ . In particular, this means that the spaces are nested: $V_0 \subset V_1 \subset \dots \subset V_L \subset V$.

Let $a : V \times V \rightarrow \mathbb{R}$ be a continuous, V -elliptic, bilinear form and $F : V \rightarrow \mathbb{R}$ be a continuous linear form and consider the problem of obtaining $u_L \in V_L$ such that

$$(15) \quad a(u_L, v) = F(v) \quad \forall v \in V_L.$$

Let $\phi_\ell^{(i)}$, $i = 1, \dots, n_\ell$, denote the usual nodal basis for V_ℓ , and ϕ_ℓ be the corresponding vector formed using the basis functions. Problem (15) is then equivalent to the matrix equation

$$\mathbf{A}_L x_L = F_L$$

where $\mathbf{A}_\ell = a(\phi_\ell, \phi_\ell)$ is the stiffness matrix, x_L is the coefficient vector of u_ℓ and $F_L = F(\phi_L)$ is the load vector. For future reference, we define the mass matrix by $\mathbf{M}_\ell = (\phi_\ell, \phi_\ell)$. The nesting of the spaces $V_\ell \subset V_{\ell+1}$ means that there exists a matrix $\mathbf{R}_{\ell+1}^\ell$ such that $\phi_\ell = \mathbf{R}_{\ell+1}^\ell \phi_{\ell+1}$, i.e. the restriction matrix with the associated prolongation matrix $\mathbf{P}_{\ell+1}^\ell$ taken to be the transpose of $\mathbf{R}_{\ell+1}^\ell$.

The first main result of the present work concerns the convergence of the Fault-Prone Two-Grid Algorithm for the special case of all $\mathcal{X}_l^{(\bullet)}$ having independent Bernoulli distributed diagonal entries with probability of fault q (which satisfies Assumption 1 with $\varepsilon = q$):

THEOREM 3. *Let $\mathcal{E}_L^{TG} = \left(\mathcal{E}_L^{S,post}\right)^{\nu_2} \mathcal{E}_L^{CG} \left(\mathcal{E}_L^{S,pre}\right)^{\nu_1}$ be the iteration matrix of the Fault-Prone Two Grid Method with Jacobi smoother and component-wise faults of rate q in prolongation, restriction, residual and in the smoother, and let \mathbf{E}_L^{TG} be its fault-free equivalent. Assume that the Assumptions 2 to 5 hold. Then*

$$(16) \quad \varrho\left(\mathcal{E}_L^{TG}\right) \leq \left\|\mathbf{E}_L^{TG}\right\|_A + C \begin{cases} q n_L^{(4-d)/2d} & d < 4, \\ q \sqrt{\log n_L} & d = 4, \\ q & d > 4, \end{cases}$$

where $\|\bullet\|_A$ is the matrix energy norm and the constant C is independent of q and L .

The proof of Theorem 3 is postponed to Appendix B. If the probability of a fault occurring is zero, i.e. $q = 0$, then the bound (16) reduces to the quantity $\left\|\mathbf{E}_L^{TG}\right\|_A$ which governs the convergence of the fault-free Two Grid Algorithm. Under the foregoing assumptions on the finite element discretization, it is known [5, 6, 18, 27] that $\left\|\mathbf{E}_L^{TG}\right\|_A \leq C_{TG} < 1$, where C_{TG} is a positive constant independent of L . If $q \in (0, 1)$, then the bound (16) depends on the failure probability q , the number of unknowns n_L and the spatial dimension d of the underlying discretization.




FIG. 2. Uniform mesh for the square domain.

Grid Method to the system arising from a discretization of the Poisson problem on a square using piecewise linears on a uniform mesh of triangles (Figure 2). In Figure 3 we plot the norm of the residual after each iteration for a range of values of failure probability q and system sizes n_L ranging from 1 thousand to 1 billion. These computations were performed on the Titan supercomputer located at Oak Ridge National Laboratory.

It is observed that when $q = 0$, the method reaches residual norm 10^{-3} in about 7 iterations uniformly for all values of n_L up to around one billion. However, as q grows, it is observed that the convergence deteriorates in all cases, with the deterioration being more and more severe as the number of unknowns is increased, until, eventually, the method fails to converge at all.

The fact that the Two Grid Method fails to converge even for such a simple classical problem means one cannot expect a more favourable behaviour for realistic practical problems.

In order to verify the scaling law suggested by (16) we examine the case $d = 2$ further. In particular, when $d = 2$, we expect convergence to fail when $q\sqrt{n_L}$ exceeds some threshold. Figure 4 shows a contour plot of the variation of the Lyapunov spectral radius (obtained by 1000 iterations of the Fault-Prone Two Grid Method) with respect to the failure probability q and the size n_L of the system. The plot indicates the boundary of the region in which the Lyapunov spectral radius exceeds unity, or equally, the region in which the fault-prone iteration no longer converges at all. It will be observed that the lines on which $q\sqrt{n_L}$ remains constant (also indicated in the plot) coincide with the contours of the plot in the region in which failure to converge occurs. This supports the scaling law suggested in the estimate (16).

4.2. Analysis of the Fault-Prone Two Grid Method with Minimal Protection. Theorem 3 indicates that the Fault-Prone Two Grid Method will generally not be resilient to faults. What remedial action, in addition to *laissez-faire*, is needed to restore the uniform convergence of the Fault-Prone Multigrid Method to that of the fault-free scheme? Theorem 4 states that if the *prolongation* operation is protected, meaning that $\mathcal{X}_L^{(P)} = \mathbf{I}$, then the rate of convergence of the Two Grid Method is independent of the number of unknowns and close to the rate of the fault-free method.

In the simplest case, in which the spatial dimension $d > 4$, the Fault-Prone Two Grid Method will converge uniformly in n_L provided that the failure probability q is sufficiently small, i.e. $q \in [0, (1 - C_{TG})/C) \subset [0, 1)$. The estimate in the cases of most practical interest, in which $d < 4$, is less encouraging with the bound depending on the size n_L of the problem being solved. In particular, the estimate means that no matter how small the probability q of a failure may be, as the size of the problem being solved grows, the bound will exceed unity suggesting that the Two Grid Method will fail to converge *at all*.

In order to illustrate this behaviour and to test the sharpness of the estimate (16), we apply the Fault-Prone Two

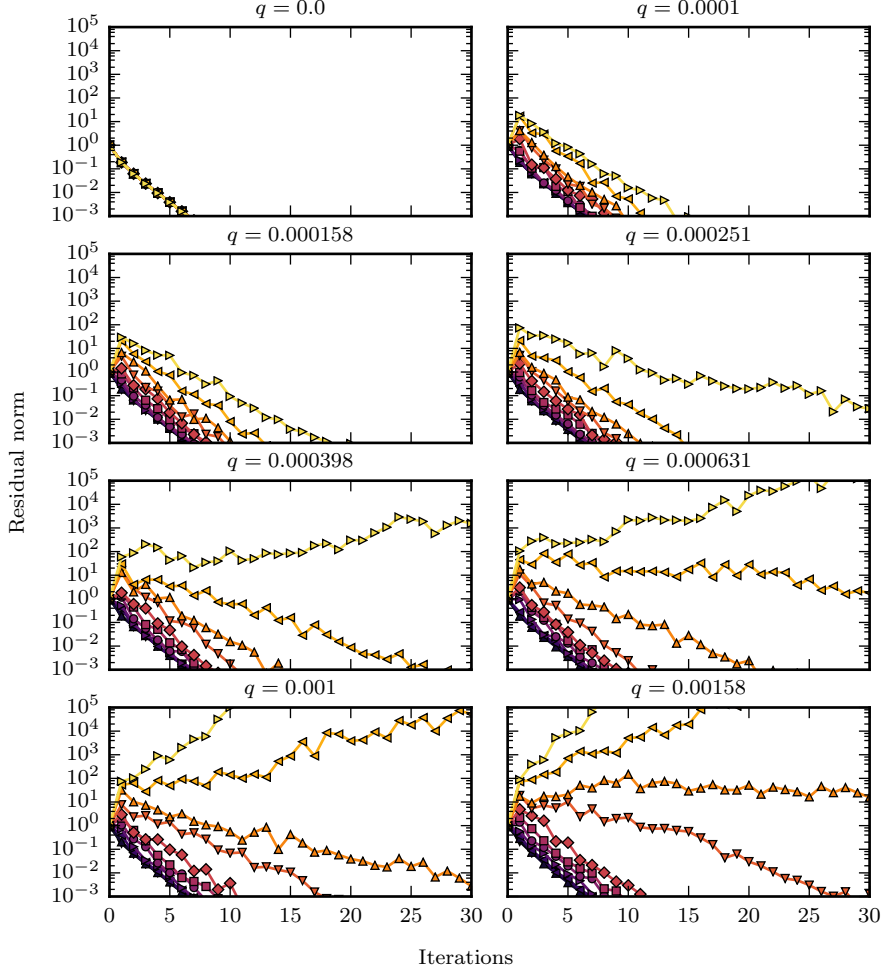



FIG. 3. Plot of the norm of the residual against iteration number for the Fault-Prone Two Grid Method in the case of discretization of the Poisson problem on a square domain.

THEOREM 4. Let $\mathcal{E}_L^{TG}(\nu_1, \nu_2) = \left(\mathcal{E}_L^{S,post}\right)^{\nu_2} \mathcal{E}_L^{CG} \left(\mathcal{E}_L^{S,pre}\right)^{\nu_1}$ be the iteration matrix of the Fault-Prone Two Grid Method with faults in smoother, residual and restriction, and protected prolongation. Provided Assumptions 1 to 5 hold, and that \mathbf{N}_L and $\mathcal{X}_L^{(S)}$ commute, we find that

$$\varrho\left(\mathcal{E}_L^{TG}(\nu_1, \nu_2)\right) \leq \min_{\substack{\mu_1 + \mu_2 = \nu_1 + \nu_2 \\ \mu_1, \mu_2 \geq 0}} \left\| \mathbf{E}_L^{TG}(\mu_1, \mu_2) \right\|_2 + C\varepsilon,$$

where the constant C is independent of ε and L .




FIG. 4. Lyapunov spectral radius $\rho(\mathcal{E}_L^{TG})$ for the iteration matrix for the Fault-Prone Two Grid Method in the case of discretization of the Poisson problem on a square domain.

The proof of Theorem 4 will be found in Appendix C. Theorem 4 makes no additional assumptions about the origin of the problem or the solver hierarchy, and allows for more general fault patterns and smootheners than Theorem 3 and as such, is more generally applicable.

We apply the Fault-Prone Two Grid Method incorporating protection of the prolongation to the same test problem as above. In Figure 6 we plot the evolution of the residual norm for varying system sizes and fault rates, and in Figure 5 we plot the Lyapunov spectral radius of the iteration matrix. Observe that the rate of convergence is independent of the number of unknowns and even stays below unity (indicating convergence) for fault probabilities as high as $q \approx 0.6$. Moreover, Figure 6 seems to indicate that while the asymptotic behaviour is independent of the number of unknowns, the first iteration is not.

In practice, protection of the prolongation might be achieved using standard techniques, such as replication or checkpointing. Since the prolongation is a local operation, its protected application could be overlapped with post-smoothing to minimize the overall performance penalty. Alternatively, some developers have proposed architectures on which certain components are less susceptible to faults. In such a case, the prolongation operator would be a candidate for this treatment.

In the present work we have proposed a simple approach to fault mitigation for iterative methods. Analysis and numerical examples showed that the Two Grid Method is not fault resilient no matter how small the error rate q may be. Moreover, we addressed the question of what is the minimal action required to restore the convergence of the algorithm. We showed that this shortcoming can be overcome by protection of the prolongation operation using standard techniques. Forthcoming work will address these issues in the case of the multigrid method.

Appendix A. Preliminaries.

Throughout the appendices, C will be a generic constant whose value can change from line to line, but which is independent of ℓ , n_ℓ , q and ε .

DEFINITION 5. Let $\mathbf{Z}, \mathbf{Y} \in \mathbb{R}^{n \times m}$, $n, m \in \mathbb{N}$. The Kronecker product $\mathbf{Z} \otimes \mathbf{Y} \in \mathbb{R}^{n^2 \times m^2}$ and the elementwise or Hadamard product $\mathbf{Z} \circ \mathbf{Y} \in \mathbb{R}^{n \times m}$ of \mathbf{Z} and \mathbf{Y} are




FIG. 5. Lyapunov spectral radius $\varrho(\mathcal{E}_L^{TG})$ for the iteration matrix for the Fault-Prone Two Grid Method with protected prolongation in the case of discretization of Poisson problem on a square domain.

defined as

$$(\mathbf{Z} \otimes \mathbf{Y})_{in+p, jm+q} = \mathbf{Z}_{ij} \mathbf{Y}_{pq}, \quad (\mathbf{Z} \circ \mathbf{Y})_{ij} = \mathbf{Z}_{ij} \mathbf{Y}_{ij}$$

for $1 \leq i, p \leq n$, $1 \leq j, q \leq m$, and the k -th Kronecker power $\mathbf{Z}^{\otimes k} \in \mathbb{R}^{n^k \times m^k}$ and the k -th Hadamard power $\mathbf{Z}^{\circ k} \in \mathbb{R}^{n \times m}$ of \mathbf{Z} as

$$\mathbf{Z}^{\otimes k} := \underbrace{\mathbf{Z} \otimes \mathbf{Z} \otimes \cdots \otimes \mathbf{Z}}_{k \text{ times}}, \quad \mathbf{Z}^{\circ k} := \underbrace{\mathbf{Z} \circ \mathbf{Z} \circ \cdots \circ \mathbf{Z}}_{k \text{ times}}.$$

THEOREM 1 (Replica trick [12]). Let \mathcal{E} be a random square matrix. Then

$$\varrho(\mathcal{E}) \leq \sqrt{\rho(\mathbb{E}[\mathcal{E} \otimes \mathcal{E}])},$$

where $\rho(\bullet)$ denotes the usual spectral radius.

Proof. Let \mathbf{E} be a non-random square matrix, then

$$\begin{aligned} \|\mathbf{E}\|_F^2 &= \text{tr}(\mathbf{E}\mathbf{E}^T) = \text{vec}(\mathbf{I}) \cdot \text{vec}(\mathbf{E}\mathbf{E}^T) \\ &= \text{vec}(\mathbf{I}) \cdot \text{vec}(\mathbf{E}\mathbf{I}\mathbf{E}^T) = \text{vec}(\mathbf{I}) \cdot \mathbf{E}^{\otimes 2} \text{vec}(\mathbf{I}), \end{aligned}$$

where $\|\bullet\|_F$ denotes the Frobenius norm and $\text{vec}(\bullet)$ denotes the vectorization of a matrix by stacking its columns. Hence, if $\lambda_{j, \otimes 2}$ and $\vec{v}_{j, \otimes 2}$ denote the eigenvalues and eigenvectors of $\mathbb{E}[\mathcal{E}^{\otimes 2}]$, sorted in descending order with respect to their absolute value, we have

$$\begin{aligned} \mathbb{E} \left[\left\| \mathcal{E}^N \right\|_F^2 \right] &= \text{vec}(\mathbf{I}) \cdot \mathbb{E} \left[(\mathcal{E}^{\otimes 2})^N \right] \text{vec}(\mathbf{I}) \\ &= \text{vec}(\mathbf{I}) \cdot \mathbb{E} [\mathcal{E}^{\otimes 2}]^N \text{vec}(\mathbf{I}) = [\text{vec}(\mathbf{I}) \cdot \vec{v}_{1, \otimes 2}]^2 \lambda_{1, \otimes 2}^N + \mathcal{O}(\lambda_{2, \otimes 2}^N). \end{aligned}$$





FIG. 6. Plot of the norm of the residual against iteration number for the Fault-Prone Two Grid Method with protected prolongation in the case of discretization of Poisson problem on square domain.

Assume without loss of generality that $|\lambda_{1,\otimes 2}| > |\lambda_{2,\otimes 2}|$. By Jensen's inequality, we find

$$\begin{aligned}
 \varrho(\mathcal{E}) &= \lim_{N \rightarrow \infty} \exp \left(\frac{1}{N} \mathbb{E} \left[\log \left\| \mathcal{E}^N \right\|_F \right] \right) \leq \lim_{N \rightarrow \infty} \exp \left(\frac{1}{2N} \log \mathbb{E} \left[\left\| \mathcal{E}^N \right\|_F^2 \right] \right) \\
 &= \sqrt{\lambda_{1,\otimes 2}} = \sqrt{\rho(\mathbb{E}[\mathcal{E}^{\otimes 2}])}
 \end{aligned}$$

and the result follows. \square

The following results are standard, but recorded here for convenience:

LEMMA 6 (Horn and Johnson [22]). *Let $\mathbf{H} = \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^T & \mathbf{C} \end{pmatrix} \in \mathbb{R}^{(p+q) \times (p+q)}$ be symmetric, $\mathbf{A} \in \mathbb{R}^{p \times p}$, $\mathbf{B} \in \mathbb{R}^{p \times q}$, $\mathbf{C} \in \mathbb{R}^{q \times q}$. Then \mathbf{H} is positive (semi-) definite if and only if \mathbf{A} and \mathbf{C} are positive (semi-) definite and there exists a contraction $\mathbf{X} \in \mathbb{R}^{p \times q}$ (i.e. all singular values are smaller or equal than 1) such that $\mathbf{B} = \mathbf{A}^{\frac{1}{2}} \mathbf{X} \mathbf{C}^{\frac{1}{2}}$.*

LEMMA 7 (Horn and Johnson [21]). *Let \mathbf{Z} , \mathbf{Y} be symmetric positive (semi-) definite matrices. Then $\mathbf{Z} \circ \mathbf{Y}$ is symmetric positive (semi-) definite.*

We now prove the following estimates for the Hadamard power of matrices:

LEMMA 8. *Let $\mathbf{Z} \in \mathbb{R}^{n \times m}$. Then*

$$\|\mathbf{Z}^{\circ 2}\|_2 \leq \max_{i=1, \dots, m} \left\| \mathbf{Z} \vec{e}_m^{(i)} \right\|_2 \max_{j=1, \dots, n} \left\| \mathbf{Z}^T \vec{e}_n^{(j)} \right\|_2 \leq \|\mathbf{Z}\|_2^2,$$

where $\{\vec{e}_m^{(i)}\}$ and $\{\vec{e}_n^{(j)}\}$ are the canonical unit basis vectors in \mathbb{R}^m and \mathbb{R}^n respectively.

Proof. Both the matrix $\begin{pmatrix} \mathbf{Z} & \mathbf{I}_n \end{pmatrix}^T \begin{pmatrix} \mathbf{Z} & \mathbf{I}_n \end{pmatrix}$ and $\begin{pmatrix} \mathbf{I}_m & \mathbf{Z}^T \end{pmatrix}^T \begin{pmatrix} \mathbf{I}_m & \mathbf{Z}^T \end{pmatrix}$ are symmetric positive semi-definite. Using Lemma 7, their Hadamard product

$$\begin{pmatrix} (\mathbf{Z}^T \mathbf{Z}) \circ \mathbf{I}_m & (\mathbf{Z}^T)^{\circ 2} \\ \mathbf{Z}^{\circ 2} & \mathbf{I}_n \circ (\mathbf{Z} \mathbf{Z}^T) \end{pmatrix}$$

is also positive semi-definite. Hence, by Lemma 6,

$$\|\mathbf{Z}^{\circ 2}\|_2^2 \leq \left\| (\mathbf{Z} \mathbf{Z}^T) \circ \mathbf{I}_n \right\|_2 \left\| (\mathbf{Z}^T \mathbf{Z}) \circ \mathbf{I}_m \right\|_2 = \max_i \left\| \mathbf{Z} \vec{e}_m^{(i)} \right\|_2^2 \max_i \left\| \mathbf{Z}^T \vec{e}_n^{(j)} \right\|_2^2.$$

The second inequality follows trivially. \square

DEFINITION 9 (Energy norms). *For matrices $\mathbf{Z} \in \mathbb{R}^{n_\ell \times n_\ell}$, we define the usual matrix energy norm $\|\mathbf{Z}\|_A$ as well as the double energy norm $\|\mathbf{Z}\|_{A^2}$ to be $\|\mathbf{Z}\|_A = \left\| \mathbf{A}_\ell^{\frac{1}{2}} \mathbf{Z} \mathbf{A}_\ell^{-\frac{1}{2}} \right\|_2$, and $\|\mathbf{Z}\|_{A^2} = \left\| \mathbf{A}_\ell \mathbf{Z} \mathbf{A}_\ell^{-1} \right\|_2$. For matrices $\mathbf{W} \in \mathbb{R}^{n_\ell^2 \times n_\ell^2}$, we define the tensor energy norm $\|\mathbf{W}\|_A$ and the tensor double energy norm $\|\mathbf{W}\|_{A^2}$ to be $\|\mathbf{W}\|_A = \left\| \left(\mathbf{A}_\ell^{\frac{1}{2}} \right)^{\otimes 2} \mathbf{Z} \left(\mathbf{A}_\ell^{-\frac{1}{2}} \right)^{\otimes 2} \right\|_2$, and $\|\mathbf{W}\|_{A^2} = \left\| \mathbf{A}_\ell^{\otimes 2} \mathbf{Z} \left(\mathbf{A}_\ell^{-1} \right)^{\otimes 2} \right\|_2$. In all cases $\|\bullet\|_2$ is the spectral norm.*

The following Lemma permits us to expand the second moment of a fault-prone iteration matrix in terms of expectations and variances of the fault matrices:

LEMMA 10. *Let \mathbf{B}_i , $i = 1, \dots, 3$, be independent random matrices. Then*

$$(17) \quad \mathbb{E} \left[(\mathbf{I} - \mathbf{B}_1)^{\otimes 2} \right] = \mathbb{E} [\mathbf{I} - \mathbf{B}_1]^{\otimes 2} + \mathbb{V} [\mathbf{B}_1]$$

and

$$(18) \quad \begin{aligned} \mathbb{E} \left[(\mathbf{I} - \mathbf{B}_1 \mathbf{B}_2 \mathbf{B}_3)^{\otimes 2} \right] = & \mathbb{E} [\mathbf{I} - \mathbf{B}_1 \mathbf{B}_2 \mathbf{B}_3]^{\otimes 2} \\ & + \mathbb{V} [\mathbf{B}_1] \mathbb{E} [\mathbf{B}_2 \mathbf{B}_3]^{\otimes 2} + \mathbb{E} [\mathbf{B}_1]^{\otimes 2} \mathbb{V} [\mathbf{B}_2] \mathbb{E} [\mathbf{B}_3]^{\otimes 2} + \mathbb{E} [\mathbf{B}_1 \mathbf{B}_2]^{\otimes 2} \mathbb{V} [\mathbf{B}_3] \\ & + \mathbb{V} [\mathbf{B}_1] \mathbb{V} [\mathbf{B}_2] \mathbb{E} [\mathbf{B}_3]^{\otimes 2} + \mathbb{V} [\mathbf{B}_1] \mathbb{E} [\mathbf{B}_2]^{\otimes 2} \mathbb{V} [\mathbf{B}_3] + \mathbb{E} [\mathbf{B}_1]^{\otimes 2} \mathbb{V} [\mathbf{B}_2] \mathbb{V} [\mathbf{B}_3] \\ & + \mathbb{V} [\mathbf{B}_1] \mathbb{V} [\mathbf{B}_2] \mathbb{V} [\mathbf{B}_3]. \end{aligned}$$

Proof. Equation (17) has already been shown in (14). To obtain the identity (18) we multiply out the tensor product, complete the square to recover the first term, and then use that

$$\mathbb{E} [\mathcal{B}_j^{\otimes 2}] = \mathbb{V} [\mathcal{B}_j] + \mathbb{E} [\mathcal{B}_j]^{\otimes 2}.$$

The proof can easily be generalised to arbitrarily many random matrices [17]. \square

Since the proof on level L only involves the levels L and $L - 1$, we will drop the first subscript and replace the second one with a subscript C for the remainder of this work.

We set

$$\mathbb{E} [\mathcal{X}^{(\bullet)}] = e^{(\bullet)} \mathbf{I}, \quad \mathbb{V} [\mathcal{X}^{(\bullet)}] = \mathbf{V}^{(\bullet)}.$$

Using Lemma 10 the second moment of the fault-prone coarse grid correction and smoothener can be written as

$$(19) \quad \mathbb{E} \left[\left(\mathcal{E}^{CG} \right)^{\otimes 2} \right] = \mathbb{E} \left[\mathcal{E}^{CG} \right]^{\otimes 2} + \mathbf{C}^{(P)} + \mathbf{C}^{(R)} + \mathbf{C}^{(\rho)} \\ + \mathbf{C}^{(P,R)} + \mathbf{C}^{(P,\rho)} + \mathbf{C}^{(R,\rho)} + \mathbf{C}^{(P,R,\rho)},$$

$$(20) \quad \mathbb{E} \left[\left(\mathcal{E}^S \right)^{\otimes 2} \right] = \mathbb{E} \left[\mathcal{E}^S \right]^{\otimes 2} + \mathbf{C}^{(S)},$$

with

$$\begin{aligned} \mathbb{E} [\mathcal{E}^{CG}] &= \mathbf{E}^{CG} + \left(1 - e^{(P)} e^{(R)} e^{(\rho)} \right) \left(\mathbf{I} - \mathbf{E}^{CG} \right), \\ \mathbf{C}^{(P)} &= \left(e^{(R)} e^{(\rho)} \right)^2 \mathbf{V}^{(P)} \left(\mathbf{P} \mathbf{A}_C^{-1} \mathbf{R} \mathbf{A} \right)^{\otimes 2}, \\ \mathbf{C}^{(R)} &= \left(e^{(P)} e^{(\rho)} \right)^2 \left(\mathbf{P} \mathbf{A}_C^{-1} \right)^{\otimes 2} \mathbf{V}^{(R)} \left(\mathbf{R} \mathbf{A} \right)^{\otimes 2}, \\ \mathbf{C}^{(\rho)} &= \left(e^{(P)} e^{(R)} \right)^2 \left(\mathbf{P} \mathbf{A}_C^{-1} \mathbf{R} \right)^{\otimes 2} \mathbf{V}^{(\rho)} \mathbf{A}^{\otimes 2}, \\ \mathbf{C}^{(P,R)} &= \left(e^{(\rho)} \right)^2 \mathbf{V}^{(P)} \left(\mathbf{P} \mathbf{A}_C^{-1} \right)^{\otimes 2} \mathbf{V}^{(R)} \left(\mathbf{R} \mathbf{A}^{\otimes 2} \right)^{\otimes 2}, \\ \mathbf{C}^{(P,\rho)} &= \left(e^{(R)} \right)^2 \mathbf{V}^{(P)} \left(\mathbf{P} \mathbf{A}_C^{-1} \mathbf{R} \right)^{\otimes 2} \mathbf{V}^{(\rho)} \mathbf{A}^{\otimes 2}, \\ \mathbf{C}^{(R,\rho)} &= \left(e^{(P)} \right)^2 \left(\mathbf{P} \mathbf{A}_C^{-1} \right)^{\otimes 2} \mathbf{V}^{(R)} \mathbf{R}^{\otimes 2} \mathbf{V}^{(\rho)} \mathbf{A}^{\otimes 2}, \\ \mathbf{C}^{(P,R,\rho)} &= \mathbf{V}^{(P)} \left(\mathbf{P} \mathbf{A}_C^{-1} \right)^{\otimes 2} \mathbf{V}^{(R)} \mathbf{R}^{\otimes 2} \mathbf{V}^{(\rho)} \mathbf{A}^{\otimes 2}, \\ \mathbb{E} [\mathcal{E}^S] &= \mathbf{E}^S + e^{(S)} \left(\mathbf{I} - \mathbf{E}^S \right), \\ \mathbf{C}^{(S)} &= \mathbf{V}^{(S)} \left(\mathbf{N} \mathbf{A} \right)^{\otimes 2}. \end{aligned}$$

When the prolongation is protected, or not subject to faults, we have $e^{(P)} = 1$ and $\mathbf{V}^{(P)} = \mathbf{0}$, so that all $\mathbf{C}^{(\bullet)}$ with a superscript containing P are zero.

Appendix B. Proof of Theorem 3.

Proof. We have that

$$e^{(\bullet)} = 1 - q, \quad \mathbf{V}^{(P)} = \mathbf{V}^{(\rho)} = \mathbf{V}^{(S)} = q(1 - q) \mathbf{K}, \quad \mathbf{V}^{(R)} = q(1 - q) \mathbf{K}_C,$$

with

$$\begin{aligned}\mathbf{K} &= \text{blockdiag} \left(\vec{e}^{(i)} \otimes \left(\vec{e}^{(i)} \right)^T, i = 1, \dots, n \right), \\ \mathbf{K}_C &= \text{blockdiag} \left(\vec{e}_C^{(i)} \otimes \left(\vec{e}_C^{(i)} \right)^T, i = 1, \dots, n_C \right).\end{aligned}$$

Here, $\vec{e}^{(i)}$ and $\vec{e}_C^{(i)}$ are the canonical unit basis vectors of \mathbb{R}^n and \mathbb{R}^{n_C} respectively. Adding and subtracting $\mathbb{E} \left[\left(\boldsymbol{\varepsilon}^{TG} \right)^{\otimes 2} \right]$ from $\mathbb{E} \left[\left(\boldsymbol{\varepsilon}^{TG} \right)^{\otimes 2} \right]$, we estimate using the energy norm on the tensor space

$$\begin{aligned}\rho \left(\mathbb{E} \left[\left(\boldsymbol{\varepsilon}^{TG} \right)^{\otimes 2} \right] \right) &\leq \left\| \mathbb{E} \left[\left(\boldsymbol{\varepsilon}^{TG} \right)^{\otimes 2} \right] \right\|_A \\ &\leq \left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^{TG} \right] \right\|_A^2 + \left\| \mathbb{E} \left[\left(\left(\boldsymbol{\varepsilon}^{S, \text{post}} \right)^{\nu_2} \boldsymbol{\varepsilon}^{CG} \left(\boldsymbol{\varepsilon}^{S, \text{pre}} \right)^{\nu_1} \right)^{\otimes 2} \right] \right\|_A \\ &\quad - \left\| \mathbb{E} \left[\left(\boldsymbol{\varepsilon}^{S, \text{post}} \right)^{\nu_2} \boldsymbol{\varepsilon}^{CG} \left(\boldsymbol{\varepsilon}^{S, \text{pre}} \right)^{\nu_1} \right] \right\|_A^2.\end{aligned}$$

We then use the subadditivity of $\|\bullet\|_A$ and equations (19) and (20) to write

$$\begin{aligned}\rho \left(\mathbb{E} \left[\left(\boldsymbol{\varepsilon}^{TG} \right)^{\otimes 2} \right] \right) &\leq \left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^{TG} \right] \right\|_A^2 \\ &\quad + \left(\left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^{CG} \right] \right\|_A^2 + \left\| \mathbf{C}^{(P)} \right\|_A + \left\| \mathbf{C}^{(R)} \right\|_A + \left\| \mathbf{C}^{(\rho)} \right\|_A \right. \\ &\quad \left. + \left\| \mathbf{C}^{(P, R)} \right\|_A + \left\| \mathbf{C}^{(P, \rho)} \right\|_A + \left\| \mathbf{C}^{(R, \rho)} \right\|_A + \left\| \mathbf{C}^{(P, R, \rho)} \right\|_A \right) \\ &\quad \times \left(\left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^S \right] \right\|_A^2 + \left\| \mathbf{C}^{(S)} \right\|_A \right)^{\nu_1 + \nu_2} \\ (21) \quad &- \left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^{CG} \right] \right\|_A^2 \left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^S \right] \right\|_A^{2(\nu_1 + \nu_2)}.\end{aligned}$$

First, we estimate the terms involving only first moments of the fault matrices. Using that \mathbf{E}^{CG} is an A -orthogonal projection and that the damped Jacobi smootheners are convergent, we find

$$\begin{aligned}\left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^{CG} \right] \right\|_A &\leq \left\| \mathbf{E}^{CG} \right\|_A + \left(1 - (1 - q)^3 \right) \left\| \mathbf{I} - \mathbf{E}^{CG} \right\|_A \leq 1 + Cq, \\ \left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^S \right] \right\|_A &\leq \left\| \mathbf{E}^S \right\|_A + q \left\| \mathbf{I} - \mathbf{E}^S \right\|_A \leq 1 + Cq,\end{aligned}$$

and therefore

$$\begin{aligned}\left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^{TG} \right] \right\|_A &\leq \left\| \mathbf{E}^{TG} \right\|_A + \left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^S \right] \right\|_A^{\nu_1 + \nu_2} \left\| \mathbb{E} \left[\boldsymbol{\varepsilon}^{CG} \right] \right\|_A - \left\| \mathbf{E}^S \right\|_A^{\nu_1 + \nu_2} \left\| \mathbf{E}^{CG} \right\|_A \\ &= \left\| \mathbf{E}^{TG} \right\|_A + Cq.\end{aligned}$$

Next, we estimate all the terms $\mathbf{C}^{(\bullet)}$ involving variances of the fault matrices. We bound

$$\frac{1}{q} \left\| \mathbf{C}^{(\rho)} \right\|_A \leq \rho \left[\left(\mathbf{A}^{\frac{1}{2}} \mathbf{P} \mathbf{A}_C^{-1} \mathbf{R} \right)^{\otimes 2} \mathbf{K} \mathbf{A}^{\otimes 2} \mathbf{K} \left(\mathbf{P} \mathbf{A}_C^{-1} \mathbf{R} \mathbf{A}^{\frac{1}{2}} \right)^{\otimes 2} \right]^{\frac{1}{2}}$$

$$\begin{aligned}
&= \rho \left[(PA_C^{-1} R A P A_C^{-1} R)^{\otimes 2} K A^{\otimes 2} K \right]^{\frac{1}{2}} \\
&= \rho \left[(PA_C^{-1} R)^{\circ 2} A^{\circ 2} \right]^{\frac{1}{2}} \\
&\leq \left\| (PA_C^{-1} R)^{\circ 2} \right\|_2^{\frac{1}{2}} \left\| A^{\circ 2} \right\|_2^{\frac{1}{2}}.
\end{aligned}$$

Here, we used that $K = K^2$ and that for compatible Z

$$(KZ^{\otimes 2}K)_{in+p, jn+q} = K_{in+p, in+p} Z_{ij} Z_{pq} K_{jn+q, jn+q} = \begin{cases} (Z^{\otimes 2})_{ij} & \text{if } i = p, j = q, \\ 0 & \text{else} \end{cases}$$

for $1 \leq i, j, p, q \leq n$. Similarly, we find

$$\begin{aligned}
\frac{1}{q} \left\| C^{(P)} \right\|_A &\leq \left\| A^{\circ 2} \right\|_2^{\frac{1}{2}} \left\| (PA_C^{-1} R)^{\circ 2} \right\|_2^{\frac{1}{2}}, \\
\frac{1}{q} \left\| C^{(R)} \right\|_A &\leq \left\| A_C^{\circ 2} \right\|_2^{\frac{1}{2}} \left\| (A_C^{-1})^{\circ 2} \right\|_2^{\frac{1}{2}}, \\
\frac{1}{q^2} \left\| C^{(P, R)} \right\|_A &\leq \left\| A^{\circ 2} \right\|_2^{\frac{1}{2}} \left\| A_C^{\circ 2} \right\|_2^{\frac{1}{2}} \left\| (A_C^{-1} R)^{\circ 2} \right\|_2, \\
\frac{1}{q^2} \left\| C^{(P, \rho)} \right\|_A &\leq \left\| A^{\circ 2} \right\|_2 \left\| (PA_C^{-1} R)^{\circ 2} \right\|_2, \\
\frac{1}{q^2} \left\| C^{(R, \rho)} \right\|_A &\leq \left\| P^{\circ 2} \right\|_2 \left\| R^{\circ 2} \right\|_2 \left\| A^{\circ 2} \right\|_2^{\frac{1}{2}} \left\| (A_C^{-1})^{\circ 2} \right\|_2^{\frac{1}{2}}, \\
\frac{1}{q^3} \left\| C^{(P, R, \rho)} \right\|_A &\leq \left\| P^{\circ 2} \right\|_2^2 \left\| A^{\circ 2} \right\|_2 \left\| (A_C^{-1} R)^{\circ 2} \right\|_2, \\
\frac{1}{q} \left\| C^{(S)} \right\|_A &\leq \left\| (NA)^{\circ 2} \right\|_2.
\end{aligned}$$

In the last inequality, we used that $V^{(S)}$ and $N^{\otimes 2}$ commute. Using Lemma 8 and Assumptions 3 and 5 we find

$$\begin{aligned}
\left\| P^{\circ 2} \right\|_2 &\leq \left\| P \right\|_2^2 \leq C, \\
\left\| A^{\circ 2} \right\|_2 &\leq \left\| A \right\|_2^2, \\
\left\| A_C^{\circ 2} \right\|_2 &\leq \left\| A_C \right\|_2^2 \leq \left\| A \right\|_2^2, \\
\left\| (PA_C^{-1} R)^{\circ 2} \right\|_2 &\leq \max_i \left\| PA_C^{-1} R \bar{e}^{(i)} \right\|_2^2 \\
&\leq \left\| PA_C^{-1} R A \right\|_2^2 \max_i \left\| A^{-1} \bar{e}^{(i)} \right\|_2^2 \\
&\leq C \max_i \left\| A^{-1} \bar{e}^{(i)} \right\|_2^2, \\
\left\| (A_C^{-1} R)^{\circ 2} \right\|_2 &\leq \max_i \left\| PA_C^{-1} \bar{e}_C^{(i)} \right\|_2 \max_j \left\| A_C^{-1} R \bar{e}^{(j)} \right\|_2 \\
&\leq C \max_i \left\| A_C^{-1} \bar{e}_C^{(i)} \right\|_2 \left\| A_C^{-1} R A \right\|_2 \max_j \left\| A^{-1} \bar{e}^{(j)} \right\|_2 \\
&\leq C \max_i \left\| A_C^{-1} \bar{e}_C^{(i)} \right\|_2 \max_j \left\| A^{-1} \bar{e}^{(j)} \right\|_2, \\
\left\| (NA)^{\circ 2} \right\|_2 &\leq \left\| NA \right\|_2^2 \leq C.
\end{aligned}$$

Now, because \mathbf{A} is the finite element discretization of a second order PDE over a quasi-uniform mesh, we obtain (see Theorem 9.11 in [15]),

$$\|\mathbf{A}\|_2 \leq Ch^{d-2}.$$

Since

$$\left\| \mathbf{A} \vec{e}^{(i)} \right\|_2 \leq Ch^{-\frac{d}{2}} \|u_L\|_{L^2},$$

where $u_L \in V_L$ is given by

$$a(u_L, v) = v(\vec{x}_i), \quad \forall v \in V_L,$$

we can apply Lemma 11 to obtain bounds for $\max_j \|\mathbf{A}^{-1} \vec{e}^{(j)}\|_2$ (and equally for $\max_i \|\mathbf{A}_C^{-1} \vec{e}_C^{(i)}\|_2$).

$$\begin{aligned} \frac{1}{q} \|\mathbf{C}^{(P)}\|_A, \frac{1}{q} \|\mathbf{C}^{(R)}\|_A, \frac{1}{q} \|\mathbf{C}^{(\rho)}\|_A, \frac{1}{q^2} \|\mathbf{C}^{(R,\rho)}\|_A &\leq C \begin{cases} h^{\frac{d-4}{2}} & d < 4, \\ \sqrt{1 + |\log h|} & d = 4, \\ 1 & d > 4, \end{cases} \\ \frac{1}{q^2} \|\mathbf{C}^{(P,R)}\|_A, \frac{1}{q^2} \|\mathbf{C}^{(P,\rho)}\|_A, \frac{1}{q^3} \|\mathbf{C}^{(P,R,\rho)}\|_A &\leq C \begin{cases} h^{d-4} & d < 4, \\ (1 + |\log h|) & d = 4, \\ 1 & d > 4, \end{cases} \\ \frac{1}{q} \|\mathbf{C}^{(S)}\|_A &\leq C. \end{aligned}$$

Therefore, we obtain from (21)

$$\rho \left(\mathbb{E} \left[\left(\boldsymbol{\varepsilon}^{TG} \right)^{\otimes 2} \right] \right) \leq \|\mathbf{E}^{TG}\|_A^2 + C \begin{cases} q^2 h^{d-4} & d < 4, \\ q^2 (1 + |\log h|) & d = 4, \\ q^2 & d > 4, \end{cases}$$

and hence

$$\varrho \left(\boldsymbol{\varepsilon}^{TG} \right) \leq \|\mathbf{E}^{TG}\|_A + C \begin{cases} q n^{\frac{4-d}{2d}} & d < 4, \\ q \sqrt{\log n} & d = 4, \\ q & d > 4, \end{cases}$$

where we used that $n \approx h^{-d}$. □

In order to conclude, we need the following technical estimate:

LEMMA 11. *Let $\Omega \in C^2$ or Ω a convex polyhedron and let $u_L \in V_L$ be the unique solution of*

$$a(u_L, v) = v(\vec{x}_i), \quad \forall v \in V_L.$$

Then

$$\|u_L\|_{L^2} \leq C \begin{cases} 1 & d < 4, \\ (1 + |\log h|)^{\frac{1}{2}} & d = 4, \\ h^{2-\frac{d}{2}} & d > 4, \end{cases}$$

where C is a constant independent of h .

Proof. Let $f \in V_L$ be the unique function that corresponds to the load and write $f = \vec{F} \cdot \phi$ with \vec{F} its coefficient vector and ϕ the vector of shape functions. Then

$$\vec{e}^{(i)} = (\phi, f)_{L^2} = (\phi, \phi)_{L^2} \cdot \vec{F} = \mathbf{M}\vec{F}.$$

Since by Theorem 9.8 in [15]

$$(22) \quad Ch^d \mathbf{I} \leq \mathbf{M} \leq Ch^d \mathbf{I},$$

we have

$$(23) \quad \|f\|_{H^m} \leq Ch^{-d} \left\| \phi^{(i)} \right\|_{H^m}.$$

Now u_L is an approximation to $u \in V$ that solves

$$a(u, v) = (f, v)_{L^2}, \quad \forall v \in V.$$

Since $f \in V \subset L^2(\Omega)$, we find by the Aubin-Nitsche Lemma [15] that

$$\|u_L - u\|_{L^2} \leq Ch^2 \|f\|_{L^2}.$$

Consider the solution $u^* \in V$ to the dual problem

$$a(v, u^*) = (u, v)_{L^2}, \quad \forall v \in V.$$

By elliptic regularity [15], we have

$$\|u^*\|_{H^2} \leq C \|u\|_{L^2}.$$

Moreover,

$$(u, u)_{L^2} = a(u, u^*) = (f, u^*)_{L^2},$$

so

$$\|u\|_{L^2}^2 \leq \|f\|_{H^{-2}} \|u^*\|_{H^2} \leq C \|f\|_{H^{-2}} \|u\|_{L^2},$$

and hence $\|u\|_{L^2} \leq C \|f\|_{H^{-2}}$. Therefore, by triangle inequality and (23), we find

$$\|u_L\|_{L^2} \leq Ch^2 \|f\|_{L^2} + C \|f\|_{H^{-2}} \leq Ch^{-d} \left(h^2 \left\| \phi^{(i)} \right\|_{L^2} + \left\| \phi^{(i)} \right\|_{H^{-2}} \right).$$

Applying the estimates for $\left\| \phi^{(i)} \right\|_{H^m}$ from Theorem 4.8 in [2], we obtain

$$\|u_L\|_{L^2} \leq C \begin{cases} 1 & d < 4, \\ (1 + |\log h|)^{\frac{1}{2}} & d = 4, \\ h^{2-\frac{d}{2}} & d > 4. \end{cases} \quad \square$$

Appendix C. Proof of Theorem 4.

Proof. Adding and subtracting $\mathbb{E} \left[\mathcal{E}^{TG}(\nu_1, \nu_2) \right]^{\otimes 2}$ from $\mathbb{E} \left[\left(\mathcal{E}^{TG}(\nu_1, \nu_2) \right)^{\otimes 2} \right]$, we estimate in $\|\bullet\|_{A^2}$

$$\rho \left(\mathbb{E} \left[\left(\mathcal{E}^{TG}(\nu_1, \nu_2) \right)^{\otimes 2} \right] \right) \leq \left\| \mathbb{E} \left[\left(\mathcal{E}^{TG}(\nu_1, \nu_2) \right)^{\otimes 2} \right] \right\|_{A^2}$$

$$\begin{aligned}
&\leq \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{TG}(\nu_1, \nu_2) \right] \right\|_{A^2}^2 + \left(\left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{CG} \right] \right\|_{A^2}^2 + \left\| \mathbf{C}^{(R)} \right\|_{A^2} + \left\| \mathbf{C}^{(\rho)} \right\|_{A^2} + \left\| \mathbf{C}^{(R, \rho)} \right\|_{A^2} \right) \\
&\quad \times \left(\left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^S \right] \right\|_{A^2}^2 + \left\| \mathbf{C}^{(S)} \right\|_{A^2} \right)^{\nu_1 + \nu_2} \\
&\quad - \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{CG} \right] \right\|_{A^2}^2 \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^S \right] \right\|_{A^2}^{2(\nu_1 + \nu_2)}.
\end{aligned}$$

We then get by Assumptions 1, 3 and 5 and $(e^\bullet)^2 \leq 1 + 2C\varepsilon + C^2\varepsilon^2 \leq 1 + C\varepsilon$ that

$$\begin{aligned}
\left\| \mathbf{C}^{(R)} \right\|_{A^2} &\leq \varepsilon \left(e^{(R)} \right)^2 \left\| \mathbf{A}_C^{-1} \mathbf{R} \mathbf{A} \right\|_2^2 \left\| \mathbf{R} \right\|_2^2 \leq \varepsilon \left(e^{(R)} \right)^2 \underline{C}_p^2 \overline{C}_p^2 C_A^2 \leq C\varepsilon, \\
\left\| \mathbf{C}^{(\rho)} \right\|_{A^2} &\leq \varepsilon \left(e^{(\rho)} \right)^2 \left\| \mathbf{P} \mathbf{A}_C^{-1} \mathbf{R} \mathbf{A} \right\|_2^2 \leq \varepsilon \left(e^{(\rho)} \right)^2 C_A^2 \leq C\varepsilon, \\
\left\| \mathbf{C}^{(R, \rho)} \right\|_{A^2} &\leq \varepsilon^2 \left\| \mathbf{A}_C^{-1} \mathbf{R} \mathbf{A} \right\|_2^2 \left\| \mathbf{R} \right\|_2^2 \leq \varepsilon^2 \underline{C}_p^2 \overline{C}_p^2 C_A^2 \leq C\varepsilon^2.
\end{aligned}$$

We also estimate

$$\left\| \mathbf{C}^{(S)} \right\|_{A^2} = \left\| \mathbf{A}^{\otimes 2} \mathbf{V}^{(S)} \mathbf{N}^{\otimes 2} \right\|_2 = \left\| (\mathbf{A} \mathbf{N})^{\otimes 2} \mathbf{V}^{(S)} \right\|_2 \leq \varepsilon \left\| \mathbf{A} \mathbf{N} \right\|_2^2 \leq C\varepsilon.$$

Here, we used Assumptions 1 and 4 and that \mathbf{N} and $\boldsymbol{\mathcal{X}}^{(S)}$ commute. Moreover,

$$\begin{aligned}
\left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{CG} \right] \right\|_{A^2} &= \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{CG} \right]^T \right\|_2 = \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{CG} \right] \right\|_2 \\
&\leq \left\| \mathbf{E}^{CG} \right\|_2 + \left| 1 - e^{(R)} e^{(\rho)} \right| \left\| \mathbf{I} - \mathbf{E}^{CG} \right\|_2 \\
&\leq C_A \left(1 + \left| 1 - e^{(R)} e^{(\rho)} \right| \right) \leq C(1 + \varepsilon), \\
\left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^S \right] \right\|_{A^2} &= \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^S \right]^T \right\|_2 = \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^S \right] \right\|_2 \\
&\leq \left\| \mathbf{E}^S \right\|_2 + \left| 1 - e^{(S)} \right| \left\| \mathbf{I} - \mathbf{E}^S \right\|_2 \leq C(1 + \varepsilon), \\
\left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{TG}(\nu_1, \nu_2) \right] \right\|_{A^2} &= \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{TG}(\nu_1, \nu_2) \right]^T \right\|_2 = \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{TG}(\nu_2, \nu_1) \right] \right\|_2 \\
&= \left\| \mathbf{E}^{TG}(\nu_2, \nu_1) \right\|_2 + \left\| \mathbb{E} \left[\boldsymbol{\mathcal{E}}^{TG}(\nu_2, \nu_1) \right] \right\|_2 - \left\| \mathbf{E}^{TG}(\nu_2, \nu_1) \right\|_2 \\
&\leq \left\| \mathbf{E}^{TG}(\nu_2, \nu_1) \right\|_2 + C\varepsilon.
\end{aligned}$$

Here, we used that by Assumption 1

$$\left| 1 - e^{(\bullet)} \right| \leq C\varepsilon, \quad \left(e^{(\bullet)} \right)^2 \leq 1 + C\varepsilon, \quad \left| 1 - e^{(R)} e^{(\rho)} \right| \leq C\varepsilon.$$

Collecting all the terms, we have

$$\rho \left(\mathbb{E} \left[\left(\boldsymbol{\mathcal{E}}^{TG}(\nu_1, \nu_2) \right)^{\otimes 2} \right] \right) \leq \left\| \mathbf{E}^{TG}(\nu_2, \nu_1) \right\|_2^2 + C\varepsilon.$$

so that we finally obtain

$$\varrho \left(\boldsymbol{\mathcal{E}}^{TG}(\nu_1, \nu_2) \right) \leq \left\| \mathbf{E}^{TG}(\nu_2, \nu_1) \right\|_2 + C\varepsilon$$

We conclude by observing that the Lyapunov spectral radius, just as the ordinary spectral radius, is invariant with respect to cyclic permutations. \square

Acknowledgements. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

REFERENCES

- [1] M. AINSWORTH AND C. GLUSA, *Numerical Mathematics and Advanced Applications - ENUMATH 2015: Proceedings of ENUMATH 2015*, Springer International Publishing, 2015, ch. Multigrid at Scale?
- [2] M. AINSWORTH, W. MCLEAN, AND T. TRAN, *The conditioning of boundary element equations on locally refined meshes and preconditioning by diagonal scaling*, SIAM Journal on Numerical Analysis, 36 (1999), pp. 1901–1932.
- [3] A. AVIZIENIS, J.-C. LAPRIE, B. RANDELL, AND C. LANDWEHR, *Basic concepts and taxonomy of dependable and secure computing*, IEEE Transactions on Dependable and Secure Computing, 1 (2004), pp. 11–33.
- [4] P. BOUGEROL AND J. LACROIX, *Products of random matrices with applications to Schrödinger operators*, vol. 8 of Progress in Probability and Statistics, Birkhäuser Boston Inc., Boston, MA, 1985.
- [5] D. BRAESS, *Finite elements. Theory, fast solvers and applications in solid mechanics. Translated from German by Larry L. Schumaker. 3rd ed.*, Cambridge: Cambridge University Press, 2007, <http://dx.doi.org/10.1017/CBO9780511618635>.
- [6] J. H. BRAMBLE, *Multigrid methods*, vol. 294, CRC Press, 1993.
- [7] J. CALHOUN, L. N. OLSON, M. SNIR, AND W. D. GROPP, *Towards a more fault resilient multigrid solver*, in Proceedings of the High Performance Computing Symposium, HPC 15, San Diego, CA, USA, 2015, Society for Computer Simulation International.
- [8] F. CAPPELLO, *Fault tolerance in petascale/exascale systems: Current knowledge, challenges and research opportunities*, International Journal of High Performance Computing Applications, 23 (2009), pp. 212–226.
- [9] F. CAPPELLO, A. GEIST, B. GROPP, L. KALE, B. KRAMER, AND M. SNIR, *Toward exascale resilience*, International Journal of High Performance Computing Applications, (2009).
- [10] F. CAPPELLO, A. GEIST, W. GROPP, S. KALE, B. KRAMER, AND M. SNIR, *Toward exascale resilience: 2014 update*, Supercomputing frontiers and innovations, 1 (2014), pp. 5–28.
- [11] 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, <http://dx.doi.org/10.1145/2304576.2304590>.
- [12] A. CRISANTI, G. PALADIN, AND A. VULPIANI, *Products of random matrices*, Springer, 1993.
- [13] T. CUI, J. XU, AND C.-S. ZHANG, *An Error-Resilient Redundant Subspace Correction Method*, ArXiv e-prints, (2013), arXiv:1309.0212.
- [14] 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*, tech. report, Lawrence Livermore National Laboratory (LLNL), Livermore, CA, Feb 2014, <http://dx.doi.org/10.2172/1149042>.
- [15] A. ERN AND J.-L. GUERMOND, *Theory and Practice of Finite Elements.*, Applied Mathematical Sciences 159. New York, NY: Springer, 2004.
- [16] H. FURSTENBERG AND H. KESTEN, *Products of Random Matrices*, The Annals of Mathematical Statistics, 31 (1960), pp. 457–469.
- [17] C. GLUSA, *Multigrid and Domain Decomposition Methods in Fault-Prone Environments*, PhD thesis, Brown University, 2017. In preparation.
- [18] W. HACKBUSCH, *Multi-grid methods and applications*, vol. 4, Springer-Verlag Berlin, 1985.
- [19] W. HACKBUSCH, *Iterative solution of large sparse systems of equations*, vol. 95 of Applied Mathematical Sciences, Springer-Verlag, New York, 1994, <http://dx.doi.org/10.1007/978-1-4612-4288-8>.
- [20] T. HERAULT AND Y. ROBERT, *Fault-Tolerance Techniques for High-Performance Computing*, Springer, 2015.
- [21] R. A. HORN AND C. R. JOHNSON, *Topics in matrix analysis*, Cambridge University press, 1991.
- [22] R. A. HORN AND C. R. JOHNSON, *Matrix analysis*, Cambridge University press, 2012.
- [23] M. HUBER, B. GMEINER, U. RÜDE, AND B. WOHLMUTH, *Resilience for Multigrid Software at the Extreme Scale*, arXiv preprint arXiv:1506.06185, (2015).
- [24] T. KRÖGER AND T. PREUSSER, *Stability of the 8-tetrahedra shortest-interior-edge partitioning*

- method*, Numerische Mathematik, 109 (2008), pp. 435–457.
- [25] M. SNIR, R. W. WISNIEWSKI, J. A. ABRAHAM, S. V. 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, 28 (2014), pp. 129–173.
 - [26] M. STOYANOV AND C. WEBSTER, *Numerical Analysis of Fixed Point Algorithms in the Presence of Hardware Faults*, SIAM Journal on Scientific Computing, 37 (2015), pp. C532–C553, <http://dx.doi.org/10.1137/140991406>.
 - [27] U. TROTTEMBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press Inc., San Diego, CA, 2001. With contributions by A. Brandt, P. Oswald and K. Stüben.
 - [28] J. N. TSITSIKLIS AND V. D. BLONDEL, *The Lyapunov exponent and joint spectral radius of pairs of matrices are hard-when not impossible-to compute and to approximate*, Mathematics of Control, Signals and Systems, 10 (1997), pp. 31–40, <http://dx.doi.org/10.1007/BF01219774>.