



HAL
open science

Automatic Algorithm-Based Fault Tolerance (AABFT) of Stencil Computations

Louis Narmour, Steven Derrien, Sanjay Rajopadhye

► **To cite this version:**

Louis Narmour, Steven Derrien, Sanjay Rajopadhye. Automatic Algorithm-Based Fault Tolerance (AABFT) of Stencil Computations. Parallel Architectures and Compilation Techniques (PACT), Oct 2023, Vienna, Austria. hal-04394874v1

HAL Id: hal-04394874

<https://inria.hal.science/hal-04394874v1>

Submitted on 15 Jan 2024 (v1), last revised 16 Jan 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Automatic Algorithm-Based Fault Tolerance (AABFT) of Stencil Computations

Louis Narmour
Univ Rennes, Inria, CNRS, IRISA,
Colorado State University
France/USA
louis.narmour@irisa.fr

Steven Derrien
Univ Rennes, Inria, CNRS, IRISA
Rennes, France
steven.derrien@irisa.fr

Sanjay Rajopadhye
Colorado State University
Fort Collins, USA
sanjay.rajopadhye@colostate.edu

ABSTRACT

In this work, we study fault tolerance of transient errors, such as those occurring due to cosmic radiation or hardware component aging and degradation, using Algorithm-Based Fault Tolerance (ABFT). ABFT methods typically work by adding some additional computation in the form of invariant checksums which, by definition, should not change as the program executes. By computing and monitoring checksums, it is possible to detect errors by observing differences in the checksum values. However, this is challenging for two key reasons: (1) it requires careful manual analysis of the input program, and (2) care must be taken to subsequently carry out the checksum computations efficiently enough for it to be worth it. Prior work has shown how to apply ABFT schemes with low overhead for a variety of input programs. Here, we focus on a subclass of programs called stencil applications, an important class of computations found widely in various scientific computing domains. We propose a new compilation scheme to automatically analyze and generate the checksum computations. To the best of our knowledge, this is the first work to do such a thing in a compiler. We show that low overhead code can be easily generated and provide a preliminary evaluation of the tradeoff between performance and effectiveness.

KEYWORDS

fault-tolerance, program transformations, polyhedral compilation, stencils

ACM Reference Format:

Louis Narmour, Steven Derrien, and Sanjay Rajopadhye. 2023. Automatic Algorithm-Based Fault Tolerance (AABFT) of Stencil Computations. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (PACT '23)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Computing technology continues to move in the direction of *more* capability and *more* complexity in *less* space with *less* power. Transient silent errors that originate in the hardware and manifest as

silent data corruption pose a serious concern for software reliability. Such errors occur due to phenomena such as cosmic radiation [14, 15] and hardware component aging and degradation [16, 17]. This is a problem of both size and scale. At one end of the spectrum, as trends push the development of smaller and lower power systems, the likelihood of encountering such errors increases [7, 18]. On the other end, silent data corruption errors have even been observed in large-scale infrastructure services running at the global scale as recently as 2021 [6]. This also has implications in the realm of High-Performance Computing (HPC) [3, 21, 22] where workloads can run for weeks and even months. Waiting months for the computation to finish only to realize that the result was incorrect has a significant impact given the time wasted. The need for robust fault tolerance today is becoming more and more prevalent as computing platforms grow in capability and complexity.

One approach to enable fault tolerance involves duplication either in the software [13, 31] or in the hardware [2]. Duplication has the highest coverage, but also the highest overhead. Other approaches avoid duplication and employ compile-time analysis to augment the software with checksums to detect errors in the memory [34]. While this is less expensive, it also has lower coverage. Silent errors that happen elsewhere, inside the floating-point arithmetic units as just one example, go undetected. Detection of silent transient errors is difficult because the errors manifest *in the data* and doing so requires running some analysis *on the computed data*, which most hardware-based fault tolerance schemes ignore.

Algorithm-based fault tolerance (ABFT) [19] has been widely studied since it was first proposed in 1984 and provides a relatively cheap way to detect, and correct, such errors *in the data*. The main idea is to augment the computation with extra work in the form of invariant checksums by exploiting algebraic identities, which by definition should remain constant valued as the program executes. By comparing checksums evaluated periodically as the program executes, we can detect errors if their difference is above some threshold. ABFT schemes have been studied on a variety of different computational kernels from dense linear algebra such as Fast-Fourier Transform networks, matrix multiplication, and more recently convolutional neural networks and stencil computations [4, 5, 8, 23, 32, 33, 38].

In this work, we focus on ABFT in the context of scientific computing. Specifically, we focus on stencil applications which are commonly used to compute approximate solutions to partial differential equations modeling physical phenomena such as (electro)dynamic wave propagation and heat flow. Even among “just” stencil computations, there is a wide range of variability. The structure of a

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

PACT '23, October 21–25, 2023, Vienna, Austria

© 2023 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

particular stencil depends on the specific properties of the particular phenomenon under study, such as the rate at which heat flows through (potentially multiple and different) physical media, just to give an example. These flow rates may be the same in all physical directions, referred to as *isotropic* diffusion in the literature, or direction-dependent, *anisotropic* [26]. To complicate things even further, these rates need not be the same everywhere, they may have different *magnitudes* at different physical points in space. All of these factors influence the particular form of the input stencil program, which directly affects the analysis required to carry out ABFT.

There are several challenges that make the application of algorithm-based fault tolerance difficult:

- (1) identification of the ABFT-applicable regions of the input program (and there may be multiple independent regions)
- (2) construction of the invariant checksums
- (3) computing the checksums efficiently enough to be worth it (i.e., with low overhead)

Each of these is largely dependent on the input program and requires very careful and manual analysis. However, there is no such *ABFT-compiler* (yet) that can perform all of the above from the input program alone. In this work, we propose a new methodology to carry out this analysis at compile time, embodied as a sequence of program transformations, which can be automated thanks to polyhedral compilation techniques. To this end, we make the following contributions:

- We show how to carry out an ABFT analysis automatically in a compiler.
- We present preliminary data on the tradeoff between overhead and error detection effectiveness of the generated code.

The rest of this paper is organized as follows. In Section 2, we provide several motivating examples showing how ABFT works on stencils illustrating the challenges mentioned above. The framework and scope of our analysis are reviewed in Section 3 followed by the description of our compiler passes in Section 4 and preliminary performance data and a discussion with respect to execution time (i.e., overhead) and effectiveness (i.e., ability to reliably detect errors) in Sections 5. Finally, we discuss related work and open questions in Sections 6 and 7.

2 MOTIVATING EXAMPLES

Algorithm-based fault tolerance works fundamentally by defining invariant checksums over subsets of the domains of the computed variables and then asserting that the checksums do not change as the program evolves. For each checksum (in general, there may be multiple) this involves constructing two algebraically equivalent expressions that compute the same value, each taken from different program slices, and then asserting that their difference is zero¹.

Stencils are typically implemented as a series of weighted convolutions. The properties mentioned above manifest in the code as different weight values and potentially different convolution kernels at different points in the domain. This directly influences where it is possible to construct the invariant checksum expressions.

¹In reality, we deal with floating point arithmetic which is not associative so the difference will never truly be zero. Instead, the common practice is to assert that it is below some sufficiently small threshold.

We illustrate this variability with several examples of increasing complexity in the following sections. One quickly appreciates the difficulty of doing this manually and can see why it is desirable to leave such analysis to a polyhedral compiler.

2.1 1D Jacobi stencil with constant weights

Consider the Jacobi 1D stencil which updates an $(N + 1)$ -element array over a series of T time steps. The primary computation of each non-boundary point comes from the weighted sum of three neighboring points from the previous time step.

$$B_{t,i} = \begin{cases} A_i & \text{if } t = 0 \\ B_{t-1,i} & \text{elif } 0 < t \leq T \text{ and } i = 0 \text{ or } i = N \\ w_0 B_{t-1,i-1} + w_1 B_{t-1,i} + w_2 B_{t-1,i+1} & \text{else} \end{cases} \quad (1)$$

In stencils, deriving the invariant checksum is achieved by building two expressions such that no two points with the same time step (i.e., the t index value of $B_{t,i}$ in Equation 1) appear in both expressions.

2.1.1 Construct invariant checksums. Let us introduce a new variable, $C_{t,l,m}$, to denote the checksum at time step t defined as the sum over the window of values of $B_{t,i}$ for $l \leq i \leq m$,

$$C_{t,l,m} = \sum_{i=l}^m B_{t,i} \quad 1 \leq l \leq m \leq N - 1 \quad (2)$$

This is illustrated as the solid boxed region in Figure 1. By substituting the definition of $B_{t,i}$ where $1 \leq l \leq m \leq N - 1$ from Equation 1 into Equation 2 we obtain,

$$C_{t,l,m} = \sum_{i=l}^m (w_0 B_{t-1,i-1} + w_1 B_{t-1,i} + w_2 B_{t-1,i+1}) \quad (3)$$

Notice that the points of B , in Equation 3, can be grouped based on the combination of w_0 , w_1 , and w_2 through which they contribute, shown by the dashed boxes in Figure 1. The middle dashed rectangle represents the points at time step $t - 1$ that contribute to C_t by all three weights. After some algebra we obtain,

$$\begin{aligned} C'_{t,l,m} \equiv & (w_0)B_{t-1,l-1} + (w_0 + w_1)B_{t-1,l} \\ & + (w_0 + w_1 + w_2) \sum_{i=l+1}^{m-1} B_{t-1,i} \\ & + (w_1 + w_2)B_{t-1,m} + (w_2)B_{t-1,m+1} \end{aligned} \quad (4)$$

which we will denote as $C'_{t,l,m}$. The right-hand sides of Equations 2 and 4 both compute the *same numerical value* using *different elements* in the domain of the variable B . Equation 2 uses elements solely from time step t , and Equation 4 from time step $t - 1$.

The only way for $C_{t,l,m}$ and $C'_{t,l,m}$ to evaluate to different values is if some error were to occur between their computation. To use this, we assert that the absolute value of their difference,

$$\Delta C_{t,l,m} \equiv |C'_{t,l,m} - C_{t,l,m}|$$

is below some threshold, large enough to be distinguishable from floating-point round-off errors, and small enough to actually detect most errors. This is discussed further in Section 5.




Figure 1: Single checksum illustration. Points contributing to $C_{t,l,m}$ in Equation 2 shown in the solid boxed region and points contributing to $C'_{t,l,m}$ in Equation 4 in dashed boxed regions. Repeating this each time step at every $(m-l)$ 'th position along i leads to an overall cost of $O(TN)$.




Figure 2: Interpolated checksum illustration after $\tau = 4$ repeated substitutions. Interpolation is more efficient than Figure 1.

2.1.2 Achieve low overhead with interpolation. For any fault tolerance scheme to be feasible, the cost of implementing it needs to be sufficiently low. However, computing and comparing checksums as shown in the previous section has a significant amount of overhead. This is because the work required to compute $\Delta C_{t,l,m}$ from Equation 2 in Figure 1 is asymptotically the same as the main stencil computation. The overhead here is too high, however, it is possible to do better.

Instead of comparing expressions across *adjacent* time steps where,

$$C' = f(B_{t-1,i}) \quad (5)$$

we could compare them across *several*, say τ , time steps such that C' is a function of only the points in B at time step $t - \tau$,

$$C' = f(B_{t-\tau,i}) \quad (6)$$

To achieve this, we can repeatedly substitute B in C' . After two substitutions we will have $C' = f(B_{t-2,i})$, after three substitutions $C' = f(B_{t-3,i})$, and so on as illustrated in Figure 2. At each step, we need to account for the combinations of weights. As we show in Section 4.1.5 since the checksum accumulates points computed by a convolution, this leads to a very regular pattern, but still, this is not something one would want to do by hand.

Expressing C' as a function of $B_{t-\tau,i}$ has the same effect as what prior work [5] refers to as “interpolating” checksums solely from other checksums at previous iterations. However, we have shown the interpolation here in the opposite order (i.e., from t to $t - 1$). Regardless of the order in which it is done, this process of interpolation is critically necessary to be able to carry out the checksum computations with low overhead. The intuition behind this can be understood by comparing the complexities of the non-interpolated scheme (Figure 1) and the interpolated scheme (Figure 2). The non-interpolated version is $O(TN)$ which has the same complexity up to a constant order as the stencil computation itself. The interpolated version, on the other hand, has the complexity $O(TN/\tau)$. This tunable parameter τ directly influences the complexity of checksum computation and consequently the overhead.

2.2 Stencils with multiple time dependencies

The interpolation process as described above is straightforward when there are only stencil dependencies from t to $t - 1$, but consider the same example now with one additional dependency on $t - 2$,

$$B_{t,i} = \begin{cases} A_i & \text{if } 0 \leq t < 2 \\ B_{t-1,i} & \text{elif } 2 \leq t \leq T \text{ and } i = 0 \text{ or } i = N \\ w_0 B_{t-1,i-1} + w_1 B_{t-1,i} + w_2 B_{t-1,i+1} & \\ \quad + w_3 B_{t-2,i} & \text{else} \end{cases} \quad (7)$$

adapted from finite-difference time-domain (FDTD) acoustic wave simulation code [25].

In this example, the process of repeated substitution results in an alternate expression of the form,

$$C' = f(B_{t-\tau,i}) + g(B_{t-\tau-1,i}) \quad (8)$$

for two different functions f and g . This is because each substitution introduces terms across the previous *two* time steps. The difficulty of identifying the correct combinations of weights here is magnified. Again, this is not something that we would want to do by hand. We show how to handle this example in Section 4.1.5 but the reader is encouraged to try to work out the equivalent version of Equation 4 and perform τ repeated substitutions obtaining the expressions for f and g to fully appreciate this for themselves.

2.3 Stencils with variable weights

Not all stencil programs are directly amenable to ABFT because it may not always be possible to do the interpolation step described in Section 2.1.2 as illustrated by the following example in Equation 9. The reason it was possible to rewrite Equation 2 as shown in Equation 4 was that the weight expressions were constant at each point in space allowing them to be factored out of the middle summation term.

Consider this slightly modified piece of stencil code in Equation 9, which is identical to Equation 1 except now there may be unique weight values at each point in space.

$$B_{t,i} = \begin{cases} A_i & \text{if } t = 0 \\ B_{t-1,i} & \text{elif } 0 < t \leq T \text{ and } i = 0 \text{ or } i = N \\ w_{i-1}B_{t-1,i-1} + w_i B_{t-1,i} + w_{i+1}B_{t-1,i+1} & \text{else} \end{cases} \quad (9)$$

It is necessary to further inspect the definition of the value of the weights, w_i before proceeding. If there exist regions with constant weights, as is common in *isotropic* stencils [26] with absorbing boundaries for example, then it is possible to carry out ABFT with low overhead. For example, isotropic stencils with absorbing boundaries typically have weights with varying values near the boundaries (i.e., close to $i = 0$ and $i = N$ in our example here) but constant “in the middle”. The definition of w_i could be of the following form,

$$w_i = \begin{cases} f(i) & 0 \leq i < M \\ 0.333 & M \leq i \leq N - M \\ f(i) & N - M < i \leq N \end{cases} \quad (10)$$

where f is some arbitrary non-constant function. In this case, it is possible to do ABFT, just not on the entire domain. The solution is to recognize that the third branch of Equation 9 can be split into three branches, one for each branch in Equation 10, of which the middle branch then has the desired constant-weights form. This is straightforward for a clever human to see, but much more difficult for a classical compiler that reasons at the statement level. All of these challenges are magnified in higher dimensions and beyond the most simple symmetrical stencils implementations, this is not something that we would want to do by hand. This kind of analysis is aptly suited to polyhedral compilation which enables reasoning at the statement *instance* level.

3 POLYHEDRAL PROGRAM REPRESENTATION

The polyhedral model [1, 9, 11, 27–29] is a mathematical formalism for reasoning about a precisely defined class of computations. It provides the technology to map high-level descriptions of compute- and data-intensive programs to a range of highly parallel targets. Polyhedral “programs” are most cleanly viewed as *equations* defined over *polyhedral domains*, evaluating an *expression* at each point therein. In the context of ABFT, it enables us to reason precisely at the program statement instance level to construct the compact sets characterizing the hyper-trapezoidal regions illustrated previously in Figure 2. The Alpha language [24, 36] is a high-level equational language that separates the specification of the program from its execution plan. The semantics of an Alpha program closely follows the program’s equivalent equational representation.

The task of extracting the polyhedral representation (i.e., the set of affine recurrence equations) from a series of nested loops with affine control structure has been well studied [10]. Given a set of affine recurrence equations, one can subsequently write the equivalent Alpha program in a straightforward manner. We do not review how to do this here and for the remainder of our discussion, we assume that the input stencil program under study is itself an Alpha program.

In Figure 3, we provide an Alpha implementation of the 1D FDTD stencil from Section 2.2. This will be our working example for the remainder of the paper. Note the layout, with the following main sections of code: inputs, outputs, locals, and equations. The

```

0: affine ftdtd [T,N,M] ->{ 1<M<N and 2<T}
1:   inputs
2:   A : {[ i ] : 0<=i<=N}
3:   outputs
4:   A_out : {[ i ] : 0<=i<=N}
5:   locals
6:   B : {[ t, i ] : 0<=t<=T and 0<=i<=N}
7:   w0,w1 : {[ i ] : 0<=i<=N}
8:   let
9:     w0[i] = case {
10:      { 0<=i<M or N-M<i<=N } : f(i);
11:      { M<=i<=N-M } : 0.284;
12:    };
13:     w1[i] = case {
14:      { 0<=i<M or N-M<i<=N } : f(i);
15:      { M<=i<=N-M } : 0.148;
16:    };
17:
18:     B[t, i] = case {
19:      { 0<=t<2 } : A[i];
20:      { t>2 and (i=0 or i=N) } : B[t-1, i];
21:      { t>2 and 0<i<N } : w0[i-1]*B[t-1, i-1]
                           + w0[i]*B[t-1, i]
                           + w0[i+1]*B[t-1, i+1]
                           + w1[i]*B[t-2, i];
22:    };
23:     A_out[i] = B[T, i];
24: .

```

Figure 3: 1D FDTD Alpha program with multiple time dependencies (from Section 2.2) and regions with variable weights (Section 2.3).

domains of the program variables are represented as sets of integer points with a syntax that closely follows that of isl (the integer

set library) [35]. Local variables are used only in the context of this program (i.e., they only have *scope* in this program).

3.1 Equations

Alpha equations are of the form below where Y is a program variable, E is an expression defined over the domain \mathcal{D}_E and f is an affine function,

$$Y[f(z)] = E[z] \quad \forall z \in \mathcal{D}_E$$

The expression E is evaluated at each point z in \mathcal{D}_E and the answer is written to the location in the variable Y specified by $f(z)$. Each Alpha expression is associated with two polyhedral domains. First, the **expression domain** of E is the domain over which E is well defined and is computed *bottom-up*. Next, the **context domain** of E is the domain over which it needs to be evaluated (i.e., where each value that contributes to an answer in the output needs to be written) and is computed *top-down*. For example, the expression $B[t, i]$ in line 23 of Figure 3 has the expression domain, $[T, N] \rightarrow \{[t, i] : 0 \leq t \leq T \text{ and } 0 \leq i \leq N\}$ with $O(TN)$ points because this is the domain over which the variable B is defined. Since it only contributes to an answer in A_{out} when $t = T$, its context domain is $[T, N] \rightarrow \{[t, i] : t = T \text{ and } 0 \leq i \leq N\}$ with $O(N)$ points. The important thing to keep in mind is that Alpha expressions are defined over polyhedral domains and when we use point-wise expressions like this, we are implicitly handling subsets of these domains.

The Alpha grammar formally defines and supports many special types of expressions. We will only review the aspects minimally needed to make our discussion in this paper self-contained.

3.2 Dependence expressions

A **dependence expression** is an expression of the following form, where f is an affine function and z is a point in the expression domain of the expression E ,

$$E[f(z)]$$

and should be understood as “the expression E evaluated at the point mapped from z by f ” or equivalently as, “read the expression E at the point $f(z)$ ”. All of the variable accesses in Figure 3 (e.g., $w[i - 1]$, $B[t - 1, i + 1]$, etc.) are dependence expressions.

3.3 Restrict expressions

A **restrict expression** is an expression of the following form, where \mathcal{D} is some domain,

$$\mathcal{D} : E$$

which should be read as “the expression E restricted to the subdomain \mathcal{D} ”. The expression domain of this restrict expression is the intersection of \mathcal{D} with the expression domain of E .

3.4 Case expressions

Expressions can also be piecewise expressions defined over multiple semicolon-delimited disjoint pieces using **case expressions**, which have the following form,

$$\text{case } \{\mathcal{D}_0 : E_0; \mathcal{D}_1 : E_1; \mathcal{D}_2 : E_2; \dots\}$$

for any number of disjoint domains \mathcal{D}_i . Here, each of the pieces is itself a restrict expression.

3.5 Reduce expressions

Finally, the Alpha language also supports reduction expressions as first-class objects. Alpha **reduce expressions** generally have the following form, where Y is a program variable, E is the expression of the reduction body with the domain \mathcal{D}_E , and f_p is an affine function,

$$Y[f_p(z)] = \bigoplus E[z] \quad \forall z \in \mathcal{D}_E$$

where the expression E , evaluated at each point z in \mathcal{D}_E , is accumulated into the element of the variable Y at the location $f_p(z)$ by the \bigoplus operator. The syntax for the reduce expression that represents this is,

$$Y[f_p(z)] = \text{reduce}(\bigoplus, f_p, \mathcal{D}_E : E[z])$$

3.6 Checksums as Alpha reductions

We use reductions to represent the checksum values over the program variables in the following sections. Equation 2 can be expressed in Alpha as,

$$C[t, l, m] = \text{reduce}(+, (t, l, m, i \rightarrow t, l, m), \{ : 1 \leq i \leq m \} : B[t, i])$$

for some new variable C representing the *family* of all possible checksum *instances* at a particular time step t over a particular range of i for $l \leq i \leq m$. We need the family to fully cover the program as discussed in Section 4.2. The reduction body here is a restrict expression with a 4-dimensional expression domain over the indices t, l, m , and i . At each point in this domain, the program variable B is read at $B[t, i]$ and accumulated into the checksum instance at $C[t, l, m]$.

In this context, the objective is to transform this reduce expression into the equivalent efficient expression over the base of the corresponding trapezoidal domain in Figure 2. Alpha supports algebraic substitution and simplification of individual expressions, which makes this possible.

4 AUTOMATIC CHECKSUM DERIVATION

To facilitate the presentation we make the following assumptions about the input stencil program,

- the outer dimension on stencil variables is interpreted as the index over time
- the equations characterizing the stencil body update points at time step t from points at strictly earlier time steps (i.e., $t - k$ for some constant $k > 0$)
- the dimensionality of all stencil variables is the same

This introduces no loss in generality because it is always possible to achieve this by reindexing the program variables. For example, consider the following Gauss-Seidel stencil equation,

$$X_{t,i} = aX_{t,i-1} + bX_{t-1,i} + cX_{t-1,i+1} \quad (11)$$

with the dependency across the same time step from $[t, i]$ to $[t, i - 1]$. In such cases, we can reindex X such that all dependencies point strictly backward in time with the mapping $\{[t, i] \rightarrow [2t + i, 2i]\}$ as illustrated in Figure 4. The interested reader can notice that this is the task of finding a set of valid scheduling hyperplanes [11].

With these assumptions in mind, principally our approach involves the following three steps to automate ABFT,

- (1) constructing checksums via algebraic substitution




Figure 4: Normalization of dependencies in Equation 11 (top) such that all dependencies subsequently point strictly backward along time t (bottom).

- (2) replicating checksums over the program domain
- (3) scheduling and code generation

which are discussed in detail in the following subsections.

4.1 Step 1 - Construct invariant checksum with computer algebra

Given an input stencil program like the one provided in Figure 3, the goal of this step is to express the stencil in such a way that we can systematically build, and subsequently insert, two algebraically equivalent expressions computing the same checksum.

4.1.1 Merge variables. Some preprocessing may be required. We will call two stencil variables, X and Y , *coupled* if they belong to the same strongly connected component in the program dependency graph. For example, if the equation for X has an expression involving Y and the equation for Y in turn has some expression involving X , then we will say that X and Y are *coupled*. There could be more than two variables as well. For example, imagine that X depends on Z which depends on Y . In this case we will say that X is coupled with both Y and Z and we will refer to the set of coupled variables as $\{X, Y, Z\}$. When there exist coupled variables like this, we merge them into a single variable as follows. Let $X_i[s]$ be a dependency expression on the i 'th such coupled variable. First, create a new higher dimensional variable X' with an additional dimension. Then embed the equations for each variable X_i in X' by associating the i 'th index value in this additional dimension with X_i . This is purely a syntactic rewrite and is always legal.

Many user-written stencil codes simulating electromagnetism, for example, contain multiple stencil variables. Consider the following example, adapted from other FDTD codes in the MathWorks library [30], with separate variables for the electric (E) and magnetic

(H) fields,

$$E_{t,i} = \begin{cases} E_{t-1,i} & t > 0 \text{ and } i = 0 \\ E_{t-1,i} - a_i(H_{t-1,i} - H_{t-1,i-1}) & t > 0 \text{ and } 0 < i < N \end{cases}$$

$$H_{t,i} = \begin{cases} H_{t-1,i} - b_i(E_{t-1,i+1} - E_{t-1,i}) & t > 0 \text{ and } 0 \leq i < N \\ H_{t-1,i} & t > 0 \text{ and } i = N \end{cases}$$

which can be merged into a single variable (M) as,

$$M_{t,i,z} = \begin{cases} M_{t-1,i,z} & \dots \text{and } z = 0 \\ M_{t-1,i,z} - a_i(M_{t-1,i,z+1} - M_{t-1,i-1,z+1}) & \dots \text{and } z = 0 \\ M_{t-1,i,z} - b_i(M_{t-1,i+1,z-1} - M_{t-1,i,z-1}) & \dots \text{and } z = 1 \\ M_{t-1,i,z} & \dots \text{and } z = 1 \end{cases}$$

with the same piecewise constraints associated with the new index z . Here z is the index associated with the new dimension, and since there are two coupled variables, z only takes on two values. The original variable E is associated with $z = 0$ and H with $z = 1$.

4.1.2 Normalize Sum of Products Expressions. Expressions for each merged stencil variable, X , and its weights, w_i are transformed into expressions of the form,

$$w_1[f_1(z)] * X[f_1(z)] + w_2[f_2(z)] * X[f_2(z)] + \dots \quad (12)$$

denoting a normal Sums of Products (SoP) expression, where the left-hand side of each product is an expression for the weights and the right-hand side is the stencil variable. For example, the expression,

$$X[t-1, i] - a[i] * (X[t-1, i+1] - X[t-1, i])$$

is rewritten as the following normal SoP expression,

$$1 * X[t-1, i] + (-1 * a[i]) * X[t-1, i+1] + a[i] * X[t-1, i]$$

This is purely a preprocessing step for the next step to identify the regions with constant weights.

4.1.3 Identify convolution domains. Given two Alpha variables X and W , let us define the convolution of W with X by the operator \otimes at the point $s \in \mathcal{D}_X$ as²,

$$W \otimes X_s = \sum_{s' \in \mathcal{D}_W} W_{s'} X_{s+s'} \quad (13)$$

Given an SoP expression involving X as defined in Equation 12, we construct a definition for W . We want to identify the subdomain \mathcal{P} that reads the same value from each weight expression appearing in the SoP. Then we define the value of $W[s]$ for each $s \in \mathcal{P}$ as the corresponding constant weight subexpression. Then we rewrite the portion of the SoP in \mathcal{P} as the convolution defined in Equation 13.

As an example, consider the following equation,

$$X[i] = w * X[i-1] \quad (14)$$

for the 1D variable X with the domain \mathcal{D}_X . The binary expression $w * X[i]$ on the right-hand side is defined over \mathcal{D}_X , and the subexpression w is read at each point in this domain. The subexpression w here should be understood as the dependence expression that reads the scalar variable w at each point in the expression domain by the affine function $f : [i] \rightarrow []$. We say that *there is reuse of w*

²The symbol “ $*$ ” is commonly used elsewhere in the literature to denote convolutions, but to avoid conflating it with Alpha’s multiplication, we use \otimes for convolutions.

in context of its use because the null space of f has a non-empty intersection with the expression domain³. This notion of reuse in context is drawn from a related polyhedral optimization called simplifying reductions [12]. We leverage this to identify the reuse space common to all products in the SoP expression.

Looking again at the body of the restrict expression on line 21 of Figure 3, this reuse analysis can be run on each of the four expressions in a bottom-up fashion. The context domain $\mathcal{D}_{\text{context}}$ of this expression is $\{[t, i] : 0 \leq t \leq T \text{ and } 0 < i < N\}$, taken from its parent restrict expression. The first weights expression $w_0[i-1]$ is only constant in the subdomain $\{[t, i] : 0 \leq t \leq T \text{ and } M \leq i-1 \leq N-M\}$. All four weights expressions are simultaneously constant in the subdomain $\mathcal{P} = \{[t, i] : 0 \leq t \leq T \text{ and } M < i < N-M\}$.

From here, we split the case branch containing this SoP into two cases. In this example, we are left with the following two expressions,

```

{: t>2 and (0<i<=M or N-M<=i<N) :
  w0[i-1]*B[t-1,i-1] + w0[i]*B[t-1,i]
  + w0[i+1]*B[t-1,i+1] + w1[i]*B[t-2,i];
{: t>2 and M<i<N-M) : // the convolution
  0.284*B[t-1,i-1] + 0.284*B[t-1,i]
  + 0.284*B[t-1,i+1] + 0.148*B[t-2,i];

```

The second expression here with constant weights can be seen as the convolution $W \otimes B_s$ where $W[-1, -1]$, $W[-1, 0]$, and $W[-1, 1]$ are 0.284 and $W[-2, 0]$ is 0.148.

4.1.4 Stencil as a composed convolution. At this point, we have identified subdomains of the stencil variables that are computed as convolutions of the form,

$$X_s = W \otimes X_s \quad (15)$$

From this definition, we can substitute X_s one time on the righthand side with Equation 15 to obtain,

$$X_s = W \otimes (W \otimes X_s) \quad (16)$$

or more generally τ times leaving,

$$X_s = W \otimes (W \otimes \dots \otimes (W \otimes X_s)) \quad (17)$$

Since convolutions are, by definition, commutative and associative, this can be equivalently expressed as,

$$X_s = W^\tau \otimes X_s \quad (18)$$

where the convolution kernel first applied to itself τ times to produce a larger kernel, W^τ which is then applied to the variable X at each point s . Note that Equations 15 and 18 compute the same value. However, the subset of points read from X differs between the two, which can be seen by referring back to Equation 13.

4.1.5 Construction of the invariant checksums. For each stencil variable X defined over the domain \mathcal{P} and its corresponding composed convolution expression defined in Equation 18, we construct one checksum expression pair C and C' . Let C be defined as the following reduction over points in \mathcal{Q} , a rectangular subdomain of \mathcal{P} ,

$$C = \sum_{s \in \mathcal{Q}} X_s \quad (19)$$

³Note that we use the `is1[35]` notation to indicate that f maps onto a 0-dimensional space denoted by the empty tuple `[]`.

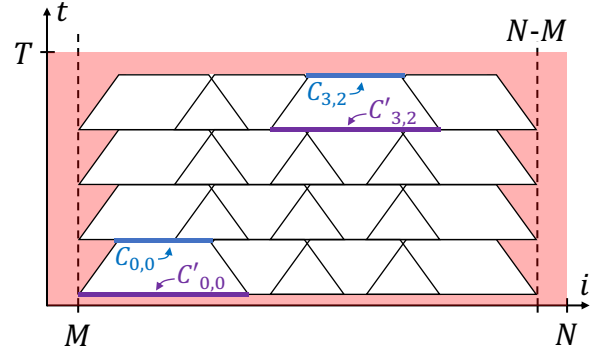


Figure 5: Striding checksum expression pairs over the corresponding convolution domain, shown here for the example in Figure 3. The subscripts $[ct, ci]$ here denote the ct 'th row (from bottom to top) and ci 'th column. The shaded red region illustrates the domain where duplication is needed.

Let C' denote an alternate checksum, which computes the same value, using Equation 18,

$$C' = \sum_{s \in \mathcal{Q}} (W^\tau \otimes X_s) \quad (20)$$

The points from X contributing to C are separated by τ timesteps from the points contributing to C' . Geometrically, these are the “top” and “bottom” faces of hyper-trapezoidal regions, as illustrated in Figure 2 for the 2D case. Silent errors occurring at any point inside this hyper-trapezoid result in different computed values for C and C' .

4.2 Step 2 - Program Error Protection Coverage

Now we have a single checksum expression pair for each convolution appearing in the normalized program. However, each pair only enables the detection of errors within their corresponding hyper-trapezoidal subdomain. In order to use this over the rest of the convolution, we fix the trapezoid to a constant (non-parametric) size from the family of all possible checksums and then stride it over the convolution domain (very similar to tiling for all intents and purposes) as shown in Figure 5.

Still, it is not possible to cover every point. Points in the red region in Figure 5 need to be handled with some other error detection scheme because they are not contained by any trapezoid. We can use duplication here, which is not a concern from an overhead perspective since this only occurs on the boundaries with an asymptotically smaller complexity. Let \mathcal{D}_B denote the domain of the stencil variable B defined in line 6 of Figure 3. Let \mathcal{D}_{ABFT} denote the domain of the union of all trapezoids. Let \mathcal{D}_{dup} denote the duplication domain and be defined as, $\mathcal{D}_{dup} = \mathcal{D}_B \setminus \mathcal{D}_{ABFT}$ where “ \setminus ” denotes set difference. Then we add another program variable, called B_{dup} , using the same equation as B but restricted to \mathcal{D}_{dup} .

4.3 Step 3 - Scheduling and code generation

At this point, the original stencil computation and the checksum expression pairs are completely decoupled. The stencil variables and checksum expression pairs need to be computed in a lock-step fashion because practical stencil implementations only use as much memory as required for two time steps' worth of computation⁴, then with “modulo 2” accessing used on the time dimension. Points in the checksums need to be accumulated into C and C' before they are overwritten at the next time step iteration. We can construct the schedule where at each time step, we first update the stencil variables, followed by any corresponding checksum variables if they exist, followed by duplication regions. Note that many time steps have no updates to any checksum variables since updates only occur at the bottom and top of the trapezoidal tiles. Alpha can be used to generate C code.

5 EVALUATION

In this section, we provide a qualitative and quantitative discussion of the effectiveness of our approach. Our primary contribution is automation so we emphasize the applicability of our analysis on non-trivial user-written stencil codes.

5.1 Qualitative Evaluation

As alluded to in Section 1, there are many factors that influence the particular form of a given stencil program, which are often more complicated than the trivial Jacobi stencils. We discuss two such stencil applications taken from GitHub and the MathWorks library and show how our approach can be applied. The first example simulates acoustic wave propagation in 2D [25]. The second example simulates electromagnetic field propagation in 3D using FDTD methods [37]. Both examples deal with non-trivial boundary conditions and are implemented in MatLab. It is worth noting that the semantics of MatLab resembles that of Alpha, in the sense that one does not need to specify the surrounding point loops and can, instead, write element-wise operations to entire data arrays with a single statement.

5.1.1 2D wave propagation code with absorbing boundaries. The first stencil application we consider implements a 2D wave propagation simulation [25]. Due to space constraints, we only highlight, with pseudo-code, the key pieces critical to our analysis. That said, this code has three parts. The corresponding Alpha implementation has equations like those shown in Figure 6. The references to N_x and N_z can be understood as size parameters of the stencil domain. Similarly, L_x and L_z are the size parameters of the subdomain where the absorbing boundary conditions hold. MatLab's element-wise operations “.” between the variables, p and $weights$, is naturally expressed in Alpha as “ $p[t-1, z, x] * weights[z, x]$ ” where the points in p at timestep $t - 1$ represent the entire p_2 array. The core challenge here is due to the fact that the expressions appearing in the stencil loop for the weights variable do not read the same value at every point, which means that the convolution can not be directly inferred. In this example, our approach will use the analysis in Section 4.1.3 to identify the interior convolution domain

⁴If there are time dependencies back to $t - k$, then k time steps' worth of memory is required of course.

shown in the bottom of Figure 6. Once this convolution domain is identified, we have all the information we need in order to apply the existing ABFT methodologies for stencils. In this sense, we can view this convolution as a variant of the simple 2D star stencil as illustrated in Figure 8.

```
// Computed variable
outputs
  p : {[T,Nz,Nx]}
let
  // Absorbing boundaries
  weights[z,x] = case {
    { z<Lz or Nz-Lz<z } : f(z);
    { x<Lx or Nx-Lx<x } : f(x);
    { Lx<=x<=Nx-Lx and Lz<=z<=Nz-Lz } : 1;
  }
  // Main stencil loop
  // p1 = p2 .* weights
  p[t,z,x] = p[t-1,z,x] * weights[z,x]

// Rewritten main stencil loop
p[t,z,x] = case {
  { z<Lz or Nz-Lz<z or x<Lx or Nx-Lx<x } : p[t-1,z,x] *
    weights[z,x];
  { Lx<=x<=Nx-Lx and Lz<=z<=Nz-Lz } : 1 * p[t-1,z,x];
}
```

Figure 6: Alpha components of the 2D wave propagation MatLab code. Given the code snippet on top, our analysis produces the code snippet at the bottom where the convolution can be identified.

5.1.2 3D FDTD electromagnetic wave simulation. The second example we consider is a significantly more complicated application simulating electromagnetic wave propagation in 3D [37]. Like the previous wave propagation example, this also has complex boundary conditions where the various weight variables are not constant in the regions close to the boundaries. These can be handled in the same way. In this example, there is additional complexity. Note that there are six stencil variables in total, the electric and magnetic field components for each of the physical spatial dimensions: E_x , E_y , E_z , H_x , H_y , and H_z . Note the intertwined dependencies across the variables (e.g., H_x reads values from E_z which in turn reads from H_y , and so on). Due to space constraints, we do not list the complete normalized version of the code, but after merging all six variables as described in Section 4.1.1, we are left with a program in the form of the bottom code snippet in Figure 7. Note that this is now in a form consumable by the analysis in Section 4.1.3 to identify the interior convolution domain. In the end, we are left with a subdomain over which we have inferred a convolution. Again, the shape of this convolution is a higher dimensional version of the star stencils illustrated in Figure 8.

5.2 Quantitative Evaluation

In this section, we report data on the overhead of the generated code produced by our analysis. Overhead refers to the execution time of the checksum-augmented program relative to the corresponding input program. We emphasize that the core contribution of this work is focused on automation and the machinery required to identify and generate code for the subdomains carrying out convolutions. Once the convolution subdomains have been identified,

```

// Computed variable
outputs
  Ez : {{T,Nx,Ny,Nz}}
let
// Main stencil loops
Hx[t,i,j,k] = case {
  D1 : Hx[t-1,i,j,k];
  D2 : ((Hx[t-1,i,j,k] + (Chzey[] * (Ey[t-1,i,j,k+1] -
    Ey[t-1,i,j,k]))) - (Chzez[] * (Ez[t-1,i,j+1,k] -
    Ez[t-1,i,j,k])));
};
// Hy, Hz, Ex, Ey equations not shown
...
Ez[t,i,j,k] = case {
  D11 : Ez[t-1,i,j,k];
  d12: ((Ez[t-1,i,j,k] + (Cezhy[i,j,k] * (Hy -
    Hy[t,i-1,j,k]))) - (Cezhx[i,j,k] * (Hx -
    Hx[t,i,j-1,k])));
};

// Merged single stencil variable
M[t,i,j,k,z] = W[i-1,j,k,0] * M[t-1,i-1,j,k,0] +
  W[i+1,j,k,1] * M[t-1,i+1,j,k,1] +
  W[i,j-1,k,2] * M[t-1,i,j-1,k,2] +
  W[i,j+1,k,3] * M[t-1,i,j+1,k,3] +
  ...;
W[i,j,k] = case {
  { : z=0 and D0' } : 1 + Cexhz[i,j,k] - Chzey[] + ...;
  { : z=1 and D1' } : 1 + Ceyhx[i,j,k] - Chyex[] + ...;
  ...;
}

```

Figure 7: Alpha components of the 3D FDTD code. Given the code snippet on top, we obtain code in the form shown at the bottom.

we can associate them with a corresponding version of the star stencil pattern illustrated in Figure 8.

All experiments were conducted on a machine running Linux with a 12th Gen Intel Core i7-12700K processor and 64Gb of memory. We report results for six different stencil codes. For several combinations of length L and height H listed in Table 1, we generated two versions of C code implementing the same stencil computation. One that only carries out the stencil computation, and one that carries out the stencil computation along with its corresponding checksum computations. In Figure 9, we report the relative performance of the checksum augmented code versus the baseline. We use L to denote the size of the base side length, and H to denote the height. In all cases, we only report data for hyper-trapezoidal shapes with square aspect ratios along the spatial dimensions, though this is not a limitation of our approach. For values of H near 16, we incur overhead near 10%. This corresponds to the results obtained from a manual implementation of a similar ABFT approach by Cavelan and Ciorba [5]. Our conclusion is that the automated code generation is competitive.

In addition to performance issues, an important concern for ABFT schemes in general is error sensitivity. Although it is an important issue, it is orthogonal to the automation process discussed in this work (we invite the interested reader to refer to the work of Cavelan and Ciorba [5] for a more in-depth discussion of this aspect). Nevertheless, for the sake of completeness, we have evaluated the sensitivity of the error detection capability on a simple 1D stencil and report our findings in Figure 10. We simulated errors with

Version	H	L	T	N
star1d1r	[2,4,8,12,16]	[50..200..50]	1000	4×10^9
star1d2r	[2,4,8,12,16]	[50..200..50]	1000	4×10^9
star2d1r	[1,2,4,8,12,16]	[40..200..40]	500	5000
star2d2r	[1,2,4,8,12,16]	[40..200..40]	500	5000
star3d1r	[1,2,4,8,12,16]	[20..100..20]	500	500
star3d2r	[1,2,4,8,12,16]	[20..100..20]	500	500

Table 1: Hyper-trapezoidal shape configuration explored for overhead experiments. The notation "[a..b..c]" denotes the inclusive range of values between "a" and "b" strided by "c".

random single bit flips on the array holding the primary stencil data in memory during execution of the program. To collect the data for Figure 10, we used a single precision (32-bit) data grid size of $N = 12000$ updated over $T = 100$ time steps. For each of the 32 bit flip positions, we ran 100 trials. If there was at least one checksum expression pair with a difference above the threshold of 10^{-5} then the run was counted as an error detection success. The fraction of total successes is reported for each checksum size (L) and bit flip position grouped by sign, exponent, upper, and lower fraction. Modeling faults as bit flips that propagate through the computation is the standard error-injection simulation approach used by related work on ABFT. As expected and also reported by Cavelan and Ciorba [5], we observe that as the size of the checksum region increases, it becomes harder to differentiate the presence of an error from floating-point round-off errors intrinsically present in the checksum computations. However, exploring this trade-off in depth is out of the scope of this work.

We have included complete working examples as a supporting artifact of the stencils and ABFT-augmented generated code used to collect the data for the overhead and fault injection experiments reported here. The reader can find the reference to the supporting code and documentation in the artifact appendix.

6 RELATED WORK

With any fault tolerance scheme, there is typically a trade-off between error coverage and overhead. Duplication and related schemes like Triple Modular Redundancy [20], for example, have the highest error coverage but are often too costly. Cavelan and Ciorba showed how to use ABFT on stencil computations, on effectively the star3d1r stencil shown in Figure 8, via interpolating checksum expressions over multiple time steps [5]. Their approach requires very careful and manual human analysis, however, and can not handle stencils with multiple time dependencies as with the example in Section 2.2. We report overhead and error detection results in agreement with their findings as a sanity check. Their work is a rediscovery of an even older work on ABFT in the context of stencils by Roy-Chowdhury et. al. [32]. Our work can be viewed as a generalization of these two approaches with a wider range of applicability.

There is other prior work that employs fault tolerance analysis at compile time. Tavarageri et. al. proposed a compiler-assisted detection of memory errors [34]. Their approach works by ensuring that any data written to memory locations remain unchanged during any of its subsequent uses. This is achieved by augmenting the input program with checksums at compile time to keep track of the




Figure 8: The core stencil pattern present in the resulting convolutions. The notation “X”d“Y”r should be read as “X” dimensional with a radius of “Y”.

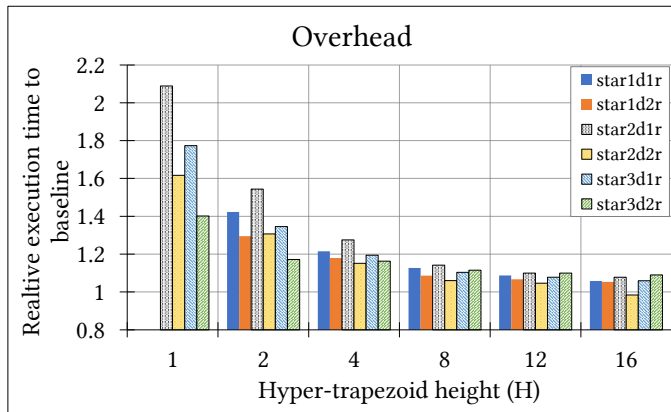


Figure 9: Shows the overhead as a function of the height H . As we expect, and was shown by prior work, larger heights lead to lower overhead.




Figure 10: Error injection experiment for the star1d1r example as a function of the base trapezoidal volume (i.e., number of points contributing to C' in Equation 20).

values *written into* and subsequently *read from* memory. Their approach has low overhead but has a lower error coverage than ABFT

approaches. Errors that occur outside of memory, in the floating point arithmetic units, for example, go undetected. Our analysis here is based on numerical properties of the computed data which means we can detect errors *anywhere* in the pipeline as long as they are not swallowed by floating-point round-off noise.

7 OPEN QUESTIONS AND FUTURE WORK

There is much more than can be explored here with respect to error detection accuracy. Our primary contribution is the automation behind the analysis and construction of ABFT checksum computations. Our approach relies on identifying regions in the stencil domain computed using convolutions. However, these weight expressions in the stencil’s convolution kernel may not be constant at each point in space or time. Stencils of this form correspond to the class of *anisotropic* partial differential equations [26]. If the weights truly are unique at each point in space, then efficient interpolation of checksums across time steps is not possible. This is because the weights can not be factored out of the expression in Equation 4. To the best of our knowledge, there has been no prior work on how to carry out ABFT efficiently on such anisotropic stencil applications and this remains an open and interesting problem.

8 CONCLUSION

We have studied the use of ABFT for stencil computations and proposed a new technique to automatically augment the input program with checksums to detect the occurrence of silent transient errors. We show that low overhead code can be easily generated thanks to polyhedral analyses and our preliminary results illustrate that there is an interesting trade-off worth exploring further.

A ARTIFACT APPENDIX

A.1 Abstract

This artifact provides a command line version of the Alpha(Z) [36] source-to-source compiler along with a set of input stencil programs that can be used to generate ABFT-hardened (checksum-augmented) C code. The provided command line scripts can then be used to compile the generated C code and to run two experiments. The first experiment validates the execution time overhead of the

ABFT-hardened program version relative to the non-hardened version showing that low overhead code can be generated. The second experiment performs error injections via bit flips and validates that checksum expression pair differences above the fault detection threshold occur can be observed as expected. All that is required to use this is a linux machine with java 11, gcc, and make.

A.2 Artifact check-list (meta-information)

- **Compilation:** AlphaZ compiler (provided as java JAR file) and make with gcc
- **Binary:** To be produced on target machine
- **Execution:** Automated via provided command line script
- **Experiments:** Execution time overhead; error injection and detection.
- **How much disk space required (approximately)?:** 300 MB
- **How much time is needed to prepare workflow (approximately)?:** <20 min (to install system dependencies if necessary)
- **How much time is needed to complete experiments (approximately)?:** 27 hours (15 for overhead experiment + 12 for error injection experiment)
- **Publicly available?:** Yes, <https://doi.org/10.5281/zenodo.8275234>
- **Code licenses (if publicly available)?:** MIT License

A.3 Description

A.3.1 How to access. Clone the following GitHub repository at <https://github.com/lnarmour/aabft-stencils-artifact>. A copy of this repository at the time of writing⁵ has been made available via archive at <https://doi.org/10.5281/zenodo.8275234>.

A.3.2 Hardware dependencies. 12 GB of memory

A.3.3 Software dependencies. java 11, gcc, make

A.4 Installation

See the README at <https://github.com/lnarmour/aabft-stencils-artifact>.

A.5 Experiment workflow

Two scripts are provided, see the README at <https://github.com/lnarmour/aabft-stencils-artifact> for usage detail.

A.6 Evaluation and expected results

See the README at <https://github.com/lnarmour/aabft-stencils-artifact>.

A.7 Methodology

Submission, reviewing and badging methodology:

- <https://www.acm.org/publications/policies/artifact-review-and-badging-current>
- <http://cTuning.org/ae/submission-20201122.html>
- <http://cTuning.org/ae/reviewing-20201122.html>

REFERENCES

- [1] P. Feautrier. 1992. Some Efficient Solutions to the Affine Scheduling Problem. Part II. Multidimensional Time. *International Journal of Parallel Programming* 21, 6 (1992), 389–420.
- [2] N. Aggarwal, P. Ranganathan, N. Jouppi, and J. Smith. 2007. Configurable isolation: Building high availability systems with commodity multi-core processors. *SIGARCH Comput. Archit. News* 35, 2 (2007), 470–481.
- [3] G. Aupy, A. Benoit, A. Cavelan, M. Fasi, Y. Robert, H. Sun, and B. Uçar. 2017. *Coping with silent errors in HPC applications*. Springer International Publishing, Cham, 269–292. https://doi.org/10.1007/978-3-319-46376-6_11
- [4] G. Bosilca, R. Delmas, J. Dongarra, and J. Langou. 2009. Algorithm-based fault tolerance applied to high performance computing. *J. Parallel and Distrib. Comput.* 69, 4 (2009), 410–416.
- [5] A. Cavelan and F. Ciorba. 2019. Algorithm-Based Fault Tolerance for Parallel Stencil Computations. In *2019 IEEE International Conference on Cluster Computing (CLUSTER)*.
- [6] H. D. Dixit, S. Pendharkar, M. Beadon, C. Mason, T. Chakravarthy, B. Muthiah, and S. Sankar. 2021. Silent Data Corruptions at Scale. *CoRR* abs/2102.11245 (2021).
- [7] R. G. Dreslinski, M. Wieckowski, D. Blaauw, D. Sylvester, and T. Mudge. 2010. Near-threshold computing: Reclaiming Moore's law through energy efficient integrated circuits. *Proc. IEEE* 98, 2 (2010), 253–266.
- [8] P. Du, A. Bouteiller, G. Bosilca, T. Herault, and J. Dongarra. 2012. 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)*. ACM, New York, NY, USA, 225–234.
- [9] P. Feautrier. 1991. Dataflow analysis of array and scalar references. *International Journal of Parallel Programming* 20, 1 (1991), 23–53.
- [10] P. Feautrier. 1991. Dataflow analysis of array and scalar references. *International Journal of Parallel Programming* 20 (1991), 23–53.
- [11] P. Feautrier. 1992. Some Efficient Solutions to the Affine Scheduling Problem. Part I. One-dimensional Time. *International Journal of Parallel Programming* 21, 5 (1992), 313–347.
- [12] Gautam and S. Rajopadhye. 2006. Simplifying reductions. In *Conference Record of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '06)*. Association for Computing Machinery, New York, NY, USA, 30–41.
- [13] M. Gupta, D. Lowell, J. Kalamatianos, S. Raasch, V. Sridharan, D. Tullsen, and R. Gupta. 2017. Compiler techniques to reduce the synchronization overhead of GPU redundant multithreading. In *2017 54th ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6. <https://doi.org/10.1145/3061639.3062212>
- [14] A. Hwang, I. Stefanovici, and B. Schroeder. 2012. Cosmic rays don't strike twice: understanding the nature of DRAM errors and the implications for system design. In *Proceedings of the Seventeenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XVII)*. ACM, New York, NY, USA, 111–122.
- [15] S. Höeffgen, S. Metzger, and M. Steffens. 2020. Investigating the effects of cosmic rays on space electronics. *Frontiers in Physics* 8 (2020).
- [16] A. Jagirdar, R. Oliveira, and T. Chakraborty. 2007. Efficient flip-flop designs for SET/SEU mitigation with tolerance to crosstalk induced signal delays. In *Proc. IEEE Workshop Silicon Errors Logic Syst. Effects*. Citeseer.
- [17] L. Juracy, M. Moreira, A. Amory, and F. Moraes. 2020. A survey of aging monitors and reconfiguration techniques. *CoRR* abs/2007.07829 (2020).
- [18] H. Kaul, M. Anders, S. Hsu, A. Agarwal, R. Krishnamurthy, and S. Borkar. 2012. Near-threshold voltage (NTV) design: Opportunities and challenges. In *Proceedings of the 49th Annual Design Automation Conference (DAC '12)*. Association for Computing Machinery, New York, NY, USA, 1153–1158.
- [19] H. Kuang-Hua and J. Abraham. 1984. Algorithm-based fault tolerance for matrix operations. *IEEE Trans. Comput.* C-33, 6 (1984), 518–528.
- [20] R. Lyons and W. Vanderkulk. 1962. The Use of Triple-Modular Redundancy to Improve Computer Reliability. *IBM Journal of Research and Development* 6, 2 (1962), 200–209.
- [21] C. Di Martino, Z. Kalbarczyk, R. Iyer, F. Baccanico, J. Fullop, and W. Kramer. 2014. Lessons learned from the analysis of system failures at petascale: the case of blue waters. In *2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. 610–621.
- [22] C. Di Martino, W. Kramer, Z. Kalbarczyk, and R. Iyer. 2015. Measuring and understanding extreme-scale application resilience: A field study of 5,000,000 HPC application runs. In *2015 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. 25–36.
- [23] T. Marty, T. Yuki, and S. Derrien. 2020. Safe overclocking for CNN accelerators through algorithm-level error detection. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 39, 12 (2020), 4777–4790.
- [24] C. Mauras. 1989. *Alpha: un langage équationnel pour la conception et la programmation d'architectures parallèles synchrones*. Ph.D. Dissertation. University of Rennes 1.
- [25] O. Ovcharenko and V. Kazei. 2018. Simple FDTD wave propagation in MATLAB. https://github.com/ovcharenko/WaveProp_in_MATLAB/tree/master/acoustic_2D_FDTD_wave_propagation
- [26] M. Patra and M. Karttunen. 2006. Stencils with isotropic discretization error for differential operators. *Numerical Methods for Partial Differential Equations* 22, 4 (2006), 936–953.

⁵At the time of writing, the top commit in this repository was b42a4d5da0adf3e34886d561c5ab3be92da8ff72

- [27] P. Quinton and V. Van Dongen. 1989. The Mapping of Linear Recurrence Equations on Regular Arrays. *Journal of VLSI Signal Processing* 1, 2 (1989), 95–113.
- [28] S. Rajopadhye, S. Purushothaman, and R. Fujimoto. 1986. On Synthesizing Systolic Arrays from Recurrence Equations with Linear Dependencies. In *Proceedings, Sixth Conference on Foundations of Software Technology and Theoretical Computer Science*. Springer Verlag, LNCS 241, New Delhi, India, 488–503.
- [29] S. V. Rajopadhye. 1989. Synthesizing Systolic Arrays with Control Signals from Recurrence Equations. *Distributed Computing* 3 (1989), 88–105.
- [30] S. Rao. 2022. 1D finite difference time domain simulation (FDTD) with perfectly matched layer (PML). <https://www.mathworks.com/matlabcentral/fileexchange/53433-1d-finite-difference-time-domain-simulation-fdtd-with-perfectly-matched-layer-pml>
- [31] G.A. Reis, J. Chang, N. Vachharajani, R. Rangan, and D.I. August. 2005. SWIFT: software implemented fault tolerance. , 243-254 pages. <https://doi.org/10.1109/CGO.2005.34>
- [32] A. Roy-Chowdhury, N. Bellas, and P. Banerjee. 1996. Algorithm-based error-detection schemes for iterative solution of partial differential equations. *IEEE Trans. Comput.* 45, 4 (1996), 394–407.
- [33] W. Syng-Jyan and N. Jha. 1994. Algorithm-based fault tolerance for FFT networks. *IEEE Trans. Comput.* 43, 7 (1994), 849–854.
- [34] S. Tavarageri, S. Krishnamoorthy, and P. Sadayappan. 2014. Compiler-assisted detection of transient memory errors (*PLDI '14*). Association for Computing Machinery, New York, NY, USA, 204–215.
- [35] S. Verdoolaege. 2010. isl: An integer set library for the polyhedral model. In *Mathematical Software – ICMS 2010*. Springer Berlin Heidelberg, Berlin, Heidelberg, 299–303.
- [36] Tomofumi Yuki, Gautam Gupta, DaeGon Kim, Tanveer Pathan, and Sanjay Rajopadhye. 2012. Alphaz: A system for design space exploration in the polyhedral model. In *International Workshop on Languages and Compilers for Parallel Computing*. Springer, 17–31.
- [37] N. Zechar. 2023. FDTD 3D and 2D Examples - ADE CPML and a Dielectric Region. <https://www.mathworks.com/matlabcentral/fileexchange/101664-fdtd-3d-and-2d-examples-ade-cpml-and-a-dielectric-region>
- [38] K. Zhao and et. al. 2021. FT-CNN: Algorithm-based fault tolerance for convolutional neural networks. *IEEE Transactions on Parallel and Distributed Systems* 32, 7 (2021), 1677–1689.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009