

FastFlip: Compositional SDC Resiliency Analysis

Keyur Joshi

kpjoshi2@illinois.edu
University of Illinois
Urbana-Champaign, USA

Rahul Singh

rahuls10@illinois.edu
University of Illinois
Urbana-Champaign, USA

Tommaso Bassetto

tommaso3@illinois.edu
University of Illinois
Urbana-Champaign, USA

Sarita Adve

sadve@illinois.edu
University of Illinois
Urbana-Champaign, USA

Darko Marinov

marinov@illinois.edu
University of Illinois
Urbana-Champaign, USA

Sasa Misailovic

misailo@illinois.edu
University of Illinois
Urbana-Champaign, USA

Abstract

To *efficiently* harden programs susceptible to Silent Data Corruptions (SDCs), developers need to invoke error injection analyses to find particularly vulnerable instructions and then selectively protect them using appropriate compiler-level SDC detection mechanisms. However, these error injection analyses are both expensive and monolithic: they must be run from scratch after even small changes to the code, such as optimizations or bug fixes. This high recurring cost keeps such software-directed resiliency analyses out of standard software engineering practices such as regression testing.

We present FastFlip, the first approach tailored to incorporate resiliency analysis seamlessly within the iterative software development workflow. FastFlip combines empirical error injection and symbolic SDC propagation analyses to enable fast and compositional error injection analysis of evolving programs. When developers modify a program, FastFlip often has to re-analyze only the modified program sections.

We analyze five benchmarks plus two modified versions per benchmark using FastFlip. FastFlip speeds up the analysis of the modified versions by 3.2× (geomean) and up to 17.2×. FastFlip selects a set of instructions to protect against SDCs that minimizes the runtime protection cost while protecting against a developer-specified target fraction of all found SDCs.

CCS Concepts: • **Software and its engineering** → **Error handling and recovery**; *Software evolution*; *Empirical software validation*.

Keywords: Error Detection, Static Analysis, Dynamic Analysis, Resiliency, Optimization.

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 the author(s) 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.
CGO'25, March 1–5, 2025, Las Vegas, United States

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 123-4-56-789012-3

<https://doi.org/https://doi.org/10.1145/3696443.3708938>

ACM Reference Format:

Keyur Joshi, Rahul Singh, Tommaso Bassetto, Sarita Adve, Darko Marinov, and Sasa Misailovic. 2025. FastFlip: Compositional SDC Resiliency Analysis. In *Proceedings of International Symposium on Code Generation and Optimization (CGO'25)*. ACM, New York, NY, USA, 15 pages. <https://doi.org/https://doi.org/10.1145/3696443.3708938>

1 Introduction

Processors are becoming increasingly susceptible to transient errors [6, 68]. The presence of Silent Data Corruptions (SDCs) in program outputs caused by such hardware-level errors during execution is difficult to detect. Software-level SDC detection techniques, such as instruction replication (e.g., [17, 48, 56]) are particularly attractive for protecting against SDCs, as they can be used on existing hardware. However, replicating all instructions leads to unacceptably high runtime overhead. To reduce runtime overhead to sustainable amounts, we can only *selectively* duplicate those instructions where errors are most likely to cause SDCs (e.g., [20, 32, 63]).

We can find vulnerable instructions using *instruction-level error injection analysis*, which injects errors one at a time into the dynamic instructions of a program during its execution, and records the effect on the output. For targeted protection, these analyses must provide *per-instruction* information on how errors in that instruction affect the output (e.g., [24, 67]), as opposed to just using sampling to provide overall statistical estimates of the program's vulnerability (e.g., [43, 53]). Such instruction-level resiliency analyses are time-consuming, requiring thousands of core-hours even for small programs.

This high analysis cost impedes the use of precise resiliency analyses in the iterative software development cycle, in which programmers regularly fix bugs, add features, and optimize their code. Each modification is frequently integrated into the code base, automatically compiled and tested to ensure the absence of bugs [69]. However, all previously proposed error injection analyses (e.g., [11, 12, 18, 30, 31, 39–41, 43, 54, 59, 66, 67]) must be re-executed on the full modified program after any (even minimal) code change, rendering them impractical.

We instead advocate for a *compositional* and *incremental* approach that partially reuses the results of error injection analyses from an old program version to reduce the analysis cost as the program evolves. We can divide the program into

sections (e.g., function calls, code blocks, or loop nests), inject errors in each section separately, and combine the per-section results to get the program’s SDC vulnerability results. When developers modify the program, we re-analyze only the sections impacted by the change and reuse the results for the sections unaffected by the modification. This is the first step toward creating a software engineering discipline for hardware errors and resilience alongside the current, well-trodden, software engineering discipline for software bugs, ensuring hardening and functional correctness are both preserved.

Designing such an approach is challenging! The approach must propagate SDCs occurring within the output of one section through downstream sections to determine the SDCs in the final output. Similarly, errors in one section can corrupt data that will be used only by subsequent sections, thus causing unexpected side effects without generating SDCs in the output of the current section. Finally, the approach should be general and support various existing resiliency analyses.

Our work. We present FastFlip, the first systematic approach for compositional and incremental error injection analysis of programs. FastFlip’s theoretical foundation describes the conditions in which combining existing instruction-level error injection analyses and symbolic error propagation analyses is possible, and its practical framework specifies how to compute the impact of injected errors on a program’s outputs. This allows FastFlip to utilize current and future advances in both sub-analyses to efficiently find vulnerable instructions.

When FastFlip first analyzes a program, FastFlip 1) *performs an error injection analysis* of each program section to find errors that cause SDCs, 2) *uses an SDC propagation analysis* to determine how SDCs propagate from one section to another to affect the final output, 3) *records the analysis results* for reuse on future program versions, and 4) *uses the analysis results to select a set of static instructions to protect* that minimizes runtime protection cost for a given target protection against SDC-causing errors. FastFlip also correctly accounts for side effects that only occur due to errors, and can adaptively adjust its results to meet SDC protection targets. When developers modify a program, FastFlip can reuse large portions of its analysis results: FastFlip only needs to rerun the expensive error injection analysis on the modified program sections and those downstream sections which receive a different input due to modified program semantics. By reusing the analysis results of other sections, FastFlip can save significant time.

For our evaluation, we instantiate FastFlip with 1) the Approxilyzer [67] per-instruction error injection analysis on top of an architectural simulator [66] and 2) the Chisel [47] SDC propagation analysis. We compare FastFlip against a baseline Approxilyzer-only error injection analysis that treats the entire program as a single section. For comparison, we use two key metrics used by previous work [23, 48, 56] on SDC protection via selective instruction duplication: 1) the *value* of protection (i.e., the coverage / fraction of SDC-causing errors detected), and 2) the *dynamic cost* of protection (e.g.,

its runtime overhead). We analyzed each benchmark both before and after making two modifications. FastFlip provides a 3.2× speedup (geomean) over Approxilyzer for analyzing the modified programs, with minimal loss in protection value or increase in cost.

Contributions. This paper makes several contributions:

- *FastFlip:* We present FastFlip, the first approach for fast, compositional SDC error injection analysis of programs. FastFlip uses error injection and SDC propagation analyses to separately analyze program sections and then combine the analysis results. FastFlip then selects a set of instructions to protect to detect a target fraction of SDC-causing errors while minimizing the runtime cost of protection.
- *Instantiation:* We realize FastFlip using the Approxilyzer error injection analysis and the Chisel SDC propagation analysis. This novel combination allows FastFlip to analyze the effect of SDC-causing bitflips in architectural registers in the dynamic instructions of a program.
- *Evaluation:* We analyze five benchmarks with FastFlip, plus two modifications per benchmark (i.e., 15 versions total). FastFlip can analyze the modified benchmarks on average 3.2× faster and up to 17.2× faster than the Approxilyzer-only approach. For all benchmark versions, FastFlip successfully protects against the target fraction of SDC-causing errors for a similar cost as the Approxilyzer-only approach.

2 Background

2.1 Error Injection Analyses

Error injection analyses first find potential error sites at various points in an error-free execution of the program. These *error sites* can be bits in various registers, caches, etc. The analysis injects errors at each site one at a time and then executes the rest of the program to record the effect of the error on the final output. Such analyses operate at different levels of abstraction, including hardware [18], assembly [67], and IR [11]. An error can have various effects on the program output:

- The error is *masked*, i.e., the program output is unaffected.
- The error causes a *crash* or other unrecoverable error leading to program termination.
- The error greatly extends the program runtime (e.g., by creating a long loop), causing a *timeout*.
- The error causes a *detectable* output change (e.g., by producing an incorrectly formatted output).
- The error causes an undetectable output change, known as a *Silent Data Corruption (SDC)*.

The analysis result maps each error site to the outcome of an error at that site. Crashes, timeouts, and detectable output changes can be handled through relatively lightweight mechanisms such as checkpoints. SDC outcomes are more insidious and require more expensive methods such as task or instruction duplication for detection. However, many applications can tolerate small errors in their outputs, such as media/signal processing and data science [61]. For such applications, it may

not be necessary to protect instructions where errors cause mostly *acceptably small* SDCs (SDC-Good) and just a few unacceptably large SDCs (SDC-Bad). Thus, similar to Approxilyzer [67] we further classify SDCs as SDC-Good or SDC-Bad based on a developer-defined and application-specific threshold ϵ . For applications that do not tolerate any SDC, ϵ is 0.

Analyses such as Approxilyzer aim to provide information on the outcome of errors at *all* error sites of a particular class within a program's execution on a specific input (e.g., [24, 67]). These instruction-level analyses are slower than approaches that do random sampling of error sites [4, 52], but can instead precisely identify vulnerable instructions that can then be protected/hardened during compilation [2, 17, 20, 29, 48, 56, 72].

2.2 SDC Propagation Analyses

SDC propagation analyses determine how SDCs within a program's input, or those introduced during execution, are propagated and amplified by the program up to the output. An SDC bound $\Delta(o) \leq f(\Delta(i))$ states that the SDC $\Delta(o)$ in the output o of a code section is at most a function f of the SDC $\Delta(i)$ in the input i . Many SDC propagation analyses, such as Chisel [47], conservatively and soundly analyze different control flow paths caused by SDCs and its impact on the output.

Sensitivity analysis [10] is a component of SDC propagation analyses that determines how sections of code amplify SDCs present within their inputs. In particular, *local* sensitivity analysis focuses on determining the effect of perturbations around a single input value. This analysis varies an input x_0 to a program section s by various amounts φ up to φ_{\max} . The analysis then executes s to calculate the output perturbation $|s(x_0 + \varphi) - s(x_0)|$ and calculates the SDC amplification factor K , which is the Lipschitz constant [13] for s at x_0 :

$$K = \max_{\varphi \leq \varphi_{\max}} \frac{|s(x_0 + \varphi) - s(x_0)|}{\varphi} \quad (1)$$

We can approximate K by sampling a set of φ values [70] or calculate its upper bound using static analysis (e.g., [13, 16, 35]).

3 Example

Lower-Upper decomposition (LUD) is a key matrix operation used in many applications. The blocked LUD algorithm consists of an outer loop with four sections that process various subsets of matrix blocks. We demonstrate FastFlip on the blocked LUD benchmark from the Splash-3 suite [57] for an example 16×16 input matrix with an 8×8 block size. Algorithm 1 shows the pseudocode of blocked LUD. In each iteration k of the loop, the loop body executes four sections s_{k1}, \dots, s_{k4} in sequence. Each section updates only one block.

Hardware errors may occasionally occur in computations using this operation. While memory can be protected using ECC, data currently being processed by the CPU is more vulnerable. If a bitflip causes an SDC, the corruption may not be detected and the user will receive a wrong answer. Here, we use the common *single-event upset* error model which is

Algorithm 1 Blocked LUD pseudocode.

```

Input  $blks$ : matrix blocks;  $n$ : blocks per dimension
Modifies  $blks$ 
1: for  $k \leftarrow 0$  to  $n-1$  do
2:   LU0( $blks[k,k]$ ) ← section  $s_{k1}$ 
3:   for  $i \leftarrow k+1$  to  $n-1$  do ] section  $s_{k2}$ 
4:     BDIV( $blks[k,i], blks[k,k]$ )
5:   for  $j \leftarrow k+1$  to  $n-1$  do ] section  $s_{k3}$ 
6:     BMODD( $blks[j,k], blks[k,k]$ )
7:   for  $i \leftarrow k+1$  to  $n-1$  do ] section  $s_{k4}$ 
8:     for  $j \leftarrow k+1$  to  $n-1$  do
9:       BMOD( $blks[j,i], blks[k,i], blks[j,k]$ )

```

widely used in previous work e.g. [24, 67]. We further assume that the error occurs in a random bit in an architectural register within a random dynamic instruction in the execution.

3.1 FastFlip Analysis

A developer can use an error injection analysis like Approxilyzer [67] (details in Section 5.1) to systematically simulate errors and determine which bitflips cause SDCs. While it gives a detailed map of vulnerable instructions, Approxilyzer requires over 600 core-hours for LU, and must be rerun from scratch after each modification to the program.

FastFlip's per-section analysis. Here, we describe how FastFlip calculates the SDC introduction and propagation characteristics (i.e., an *SDC specification*) of the 1st code section in the 2nd iteration of the LUD computation (referred to as s_{21}) given its input data I_{21} . FastFlip repeats the following process for each section s of the full program execution T :

- FastFlip uses Approxilyzer on s_{21} in isolation to determine the effect of bitflips in each instruction in s_{21} on its output $O_{s_{21}}$. Some bitflips lead to SDCs in $O_{s_{21}}$. We denote as $\varphi_{s_{21}}$ the magnitude of SDC introduced into $O_{s_{21}}$ due to such bitflips.
- In addition, s_{21} can also amplify SDCs already present within its input $I_{s_{21}}$ due to a bitflip in previous computation. FastFlip uses a local sensitivity analysis to calculate the amplification factor. This sensitivity analysis calculates that if the magnitude of SDC present in $I_{s_{21}}$ is $\Delta(I_{s_{21}})$, the resulting SDC in $O_{s_{21}}$ will be at most $f_{s_{21}}(\Delta(I_{s_{21}})) = 3.2\Delta(I_{s_{21}})$, i.e., s_{21} amplifies input SDCs by at most $3.2\times$.
- FastFlip combines these formulas to create a symbolic *SDC specification* for the section s_{21} . Under the single bitflip error model, the *total* magnitude of SDC in $O_{s_{21}}$ ($\Delta(O_{s_{21}})$) is upper-bounded by the sum of the propagated SDC ($f_{s_{21}}(\Delta(I_{s_{21}}))$) and the SDC potentially introduced by a bitflip in s_{21} ($\varphi_{s_{21}}$):

$$\Delta(O_{s_{21}}) \leq \varphi_{s_{21}} + f_{s_{21}}(\Delta(I_{s_{21}})) \quad \text{where } f_{s_{21}}(\Delta(I_{s_{21}})) = 3.2\Delta(I_{s_{21}})$$

Calculating an end-to-end SDC specification. FastFlip next provides these SDC specifications for all sections to Chisel [47], an SDC propagation analysis, plus a specification of data flow between sections (details in Section 5.1). Chisel uses this information to propagate potential SDCs caused by




Figure 1. FastFlip’s protection value (top) and protection cost (bottom), which closely match the target value and Approxilyzer’s protection cost, respectively (under 0.1% difference).

bitflips up to the final output and calculates the whole execution’s SDC specification. In this case, for two iterations of LU with four sections, Chisel calculates the following expression:

$$\Delta(O_{fin}) \leq 4174.8\varphi_{s_{11}} + 434.3\varphi_{s_{12}} + 28.8\varphi_{s_{13}} + 3.2\varphi_{s_{14}} \quad (2)$$

$$+ \varphi_{s_{21}} + \varphi_{s_{22}} + \varphi_{s_{23}} + \varphi_{s_{24}}.$$

Here $\varphi_{s_{xy}}$ is a symbolic variable representing the SDC potentially introduced into section y in iteration x . The numerical coefficient next to each $\varphi_{s_{xy}}$ represents the Chisel-calculated *total* amplification of $\varphi_{s_{xy}}$ by sections downstream of the error injection point (the coefficients depend on the program’s input matrix data). FastFlip uses Equation 2 to propagate different SDCs from each section to the final output.

Selecting instructions to protect. FastFlip adapts the value and cost model from [23] to select a set of instructions to protect. FastFlip associates each static instruction pc with 1) the *value* $v(pc)$ of protecting it, i.e., the number of SDC-causing bitflips at pc in the program execution T , and 2) the *cost* $c(pc)$ of said protection, i.e., the number of dynamic instances of pc in the program execution T .

The value and cost of protecting a set of instructions are the sum of the value and cost of protecting each instruction in the set. This creates a trade-off space of total protection value and cost corresponding to each possible subset of instructions (see Figure 1; bottom plot). Given a target total SDC protection value, FastFlip aims to select a subset of instructions that minimize the total protection cost. This is a 0–1 knapsack optimization problem, which FastFlip solves via dynamic programming.

3.2 FastFlip Results

We compare FastFlip’s results with those of an Approxilyzer-only approach that analyzes the whole program execution at once. We assume that a developer wants to protect against at least 90% of SDC-causing bitflips.

Value. The top plot in Figure 1 shows the value of protecting FastFlip’s selection of instructions against SDCs. The X and Y-Axes show the target and achieved value, respectively. The solid blue line shows FastFlip’s achieved value, which overlaps the dotted black line showing the target value. FastFlip successfully achieves the target value for the entire target range.

Cost. The bottom plot in Figure 1 compares the cost of protecting FastFlip’s and Approxilyzer’s selections of instructions against SDCs. The X-Axis shows the target value, while the Y-Axis shows the protection cost in terms of the number of dynamic instructions which must be protected. The red dashed line and solid blue line show the cost using Approxilyzer and FastFlip’s results, respectively. The two lines overlap, and the excess of cost of FastFlip over Approxilyzer is below 0.1%.

Modifications. We next perform both analyses on two modified versions of this program. The *small* modification (a few lines of code; see Section 5.5) uses a specialized version of section s_{k4} which reduces the number of bounds checks when the matrix size is a multiple of the block size (as in our input). The *large* modification replaces section s_{k1} with a lookup table. Unlike the baseline, which must inject errors in the full execution of the modified program, FastFlip only needs to inject errors in the modified program sections, saving considerable time. FastFlip’s maximum deviation from the target value is 0.1% for these modified programs, and the excess of cost of FastFlip over Approxilyzer stays below 0.3%.

Analysis time. FastFlip requires 694 core-hours to analyze the original version of the program, compared to 602 core-hours for Approxilyzer. This is because Approxilyzer can *prune* error injections by forming equivalence classes of bitflips (i.e., bitflips that lead to the same outcome) across multiple sections. However, FastFlip saves significant time when later analyzing the modified versions of the program:

- *Small* modification: FastFlip requires 80 core-hours, compared to Approxilyzer’s 625 core-hours (7.8× faster).
- *Large* modification: FastFlip requires 94 core-hours, compared to Approxilyzer’s 441 core-hours (4.7× faster).

This shows that FastFlip’s advantage is in analyzing programs as they gradually evolve, saving time with each modification.

4 The FastFlip Approach

Figure 2 visualizes the FastFlip approach. First, FastFlip performs two sub-analyses on each program section s in the full program execution T : 1) FastFlip uses an *error injection analysis*¹ to determine the effect of each injection in s and stores the outcome, and 2) FastFlip uses a *local sensitivity analysis* to obtain an SDC propagation specification for s , and converts it into a total SDC specification for s . Second, FastFlip runs an *SDC propagation analysis*¹ over T to obtain end-to-end SDC propagation specifications. Third, FastFlip calculates concrete end-to-end SDC magnitudes to find the probability of an SDC-Bad outcome associated with each static instruction. Finally,

¹Section 4.8 describes the properties of supported sub-analyses.




Figure 2. The FastFlip approach.

FastFlip selects a set of instructions to protect with SDC detection mechanisms that minimizes the cost of protection while also ensuring that the total value of the protection against SDCs is above a developer-defined threshold.

4.1 Preliminaries

Definitions. We use the following symbols:

- T : dynamic trace of full program execution.
- s : section of the full program execution (usually a function call or execution of a code block or loop nest); $s \in T$.
- J : set of all error injection sites in T .
- J_s : set of all error injection sites in s ; $J_s \subseteq J$.
- $O_s(j)$: effect of an injection j on the outputs of s calculated by the error injection analysis.
- $i_{s,0}, \dots, i_{s,m}$ and $o_{s,0}, \dots, o_{s,n}$: inputs and outputs of s .
- $i_{T,0}, \dots, i_{T,m}$ and $o_{T,0}, \dots, o_{T,n}$: inputs and outputs of T .
- $f_{s,k}, f_{T,\lambda}, f_{T,\lambda,s}$: specifications of how the program sections propagate SDCs, calculated by the local sensitivity analysis, the SDC propagation analysis, and FastFlip respectively.
- $\varphi_{s,k}, \varphi_{*,*}, \varphi_{s,*}, \varphi_{\bar{s},*}$: symbolic variables (or sets thereof) for SDCs introduced into section outputs by errors.
- $p(j)$: probability that the error occurs at error site $j \in J$.
- $PC(j)$: maps $j \in J$ to the corresponding static instruction identifier pc . $PC(J)$ denotes the set of all static instructions of interest for error injection.
- ε_λ : maximum acceptable SDC for output $o_{T,\lambda}$ of T .
- $v(pc)$: the value of protecting the static instruction at pc .
- $c(pc)$: the cost of protecting the static instruction at pc .
- pc_{prot} : static instructions selected for SDC protection.

Analysis inputs. FastFlip accepts the full program T , its partition into sections s , a specification of how data flows between sections, the probabilities $p(j)$, SDC limits ε_λ , and the protection cost function $c(\cdot)$ as inputs.

Sections are developer-identified parts of the program that perform specific tasks, like function calls, code blocks, or loop nests. Developers can obtain the dataflow specification using standard compiler analysis passes. Expert developers can also input this data manually, as we do.

Assumptions. As in previous works [24, 67], FastFlip assumes that: 1) exactly one error occurs during the execution of the full program, and 2) the program's input is SDC-free.

4.2 Error Injection Analysis of Program Sections

FastFlip runs an error injection analysis on each program section $s \in T$ to determine the effect of errors on the outputs of s . If an injected error j causes a detectable outcome (crash, timeout, misformatted output, etc.), then the outcome $O_s(j) = detected$. Otherwise, the outcome $O_s(j) = (r_0, r_1, \dots, r_n)$, where r_k is the magnitude of SDC caused by the injection j in output $o_{s,k}$ of s . If the injection is masked for an output $o_{s,k}$, then $r_k = 0$. Depending on the application and analysis, SDC magnitude can be measured as absolute error, relative error, PSNR, etc.

4.3 SDC Propagation Analysis in Program Sections

FastFlip performs a local sensitivity analysis on each program section $s \in T$ to calculate how it amplifies SDCs present within its input. The local sensitivity analysis produces an SDC propagation specification for s of the general form:

$$\bigwedge_{k=0}^n \Delta(o_{s,k}) \leq f_{s,k}(\Delta(i_{s,0}), \dots, \Delta(i_{s,m}))$$

The specification bounds the SDC $\Delta(o_{s,k})$ in each output $o_{s,k}$ of s using a function $f_{s,k}$ of the SDC bounds of the inputs of s . FastFlip adds symbolic variables $\varphi_{s,k}$ to represent the magnitude of SDC introduced to $o_{s,k}$ during the execution of s as a result of an error within s . Under the single error model, if the input to s already contains SDC, then the error occurred in a previous program section, hence s cannot introduce additional SDC. Thus, we can write the total SDC magnitude in the outputs of s as the sum of the SDC magnitude due to an error in s and the SDC propagated by s from its input to its output:

$$\bigwedge_k \Delta(o_{s,k}) \leq f_{s,k}(\Delta(i_{s,0}), \dots, \Delta(i_{s,m})) + \varphi_{s,k} \quad (3)$$

4.4 End-to-End SDC Propagation Analysis

FastFlip runs an SDC propagation analysis on the full program T . FastFlip provides the analysis with the total SDC specifications from Equation 3 for each $s \in T$. The analysis also uses the developer-provided dataflow specification indicating how outputs of one section flow into the inputs of subsequent sections. With this information, the SDC propagation analysis calculates the relationship between errors that occur anywhere in T to the SDC in the outputs of T . It creates

an end-to-end SDC propagation specification of the form:

$$\bigwedge_{\lambda=0}^n \Delta(o_{T,\lambda}) \leq f_{T,\lambda}(\Delta(i_{T,0}), \dots, \Delta(i_{T,m}), \varphi_{*,*})$$

where $\varphi_{*,*}$ is the list of all $\varphi_{s,k}$ variables from Equation 3 across all sections. Like previous analyses [66, 67], FastFlip assumes that the input to the first section is SDC-free in order to focus on SDCs caused by the analyzed program only. So, we can simplify $f_{T,\lambda}$ by removing all $\Delta(i_{T,*})$:

$$\bigwedge_{\lambda} \Delta(o_{T,\lambda}) \leq f_{T,\lambda}(\varphi_{*,*})$$

We next create specialized versions of $f_{T,\lambda}$ by noting that, under the single error model, the φ variables for only one section can be nonzero at a time: $f_{T,\lambda,s}(\varphi_{s,*}) = f_{T,\lambda}(\varphi_{s,*}, \varphi_{s,*} = 0)$. We rewrite the end-to-end SDC propagation specification as:

$$j \in J_s \Rightarrow \bigwedge_{\lambda} \Delta(o_{T,\lambda}) \leq f_{T,\lambda,s}(\varphi_{s,*}) \quad (4)$$

Equation 4 states that, if an error occurs in section s ($j \in J_s$), then the upper bound on the SDC in output $o_{T,\lambda}$ of T is given by $f_{T,\lambda,s}(\varphi_{s,*})$, a function of the SDCs in the outputs of s .

4.5 Calculating Value of Protecting Static Instructions

FastFlip uses the injection outcomes (Section 4.2) and Equation 4 to answer the following question: *For a given static instruction identified by its program counter pc in the full execution T , what is the total probability that error injections in pc will result in SDC-Bad ($|SDC| > \epsilon_\lambda$) for any output $o_{T,\lambda}$ of T ? This is the value $v(pc)$ of protecting pc .*

Algorithm 2 Find the value of protecting static instructions.

Input • $T, J_s, PC(j), \epsilon_\lambda, p(j)$: defined in Section 4.1;
 • $O_s(j)$: outcome of injection at $j \in J_s$;
 • $f_{T,\lambda,s}$: SDC propagation specifications from Equation 4

Returns $\forall pc. v(pc)$: value of protecting pc

```

1:  $v \leftarrow \{\forall pc. pc \mapsto 0\}$ 
2: for  $s$  in  $T$  and  $j$  in  $J_s$  do
3:    $pc \leftarrow PC(j)$ 
4:   if  $O_s(j) \neq \text{detected}$  then
5:     if  $\exists \lambda. f_{T,\lambda,s}(O_s(j)) > \epsilon_\lambda$  then
6:        $v(pc) \leftarrow v(pc) + p(j)$ 
7:  $\forall pc. v(pc) \leftarrow v(pc) / \sum_{pc} v(pc)$ 

```

Algorithm 2 shows how FastFlip calculates $v(pc)$. For each error injection in each section, FastFlip checks if the error results in a detectable outcome. If not, FastFlip calculates the RHS of Equation 4 to use as an upper bound on the magnitude of SDC in the outputs of T as a result of the error (i.e., the LHS of Equation 4). If the SDC in any output is SDC-Bad, FastFlip adds the probability of that error to the value of protecting pc . Lastly, FastFlip rescales the values so that the total value of protecting all static instructions is 1.

4.6 Finding an Optimal Set of Instructions to Protect

FastFlip uses the values $v(pc)$ calculated by Algorithm 2 for each pc and the corresponding protection costs $c(pc)$ as inputs to a 0–1 knapsack optimization problem. We model the value and cost of protecting a set of instructions as the sum of the value and cost of protecting each instruction in the set. Given a developer-defined target total protection value v_{trgt} , FastFlip solves the knapsack problem via the standard dynamic programming approach to select a set of static instructions pc_{prot} to protect that minimizes the total protection cost:

$$\underset{pc_{prot} \subseteq PC(J)}{\operatorname{argmin}} \sum_{pc \in pc_{prot}} c(pc) \quad \text{such that} \quad \sum_{pc \in pc_{prot}} v(pc) \geq v_{trgt}$$

We represent the set of all static instructions of interest $PC(J)$ as a binary vector, with one bit per static instruction. A bit in the vector is set if and only if the corresponding static instruction is in pc_{prot} . Under these conditions, the objective and constraints become linear functions of binary variables.

To explore the tradeoff space between value and cost, FastFlip selects the optimal pc_{prot} for a *range* of v_{trgt} values (e.g., $v_{trgt} \in [0.9, 1.0]$). This process corresponds to solving the value / cost multi-objective optimization problem using the ϵ -constraint method [46] (i.e., turning one of the objectives into a constraint) to obtain Pareto-optimal choices for pc_{prot} .

4.7 Composability

When developers modify a program section, FastFlip must rerun the error injection and local sensitivity analysis on the modified program section. If the modification changes the input to a downstream section by changing the modified section's semantics, FastFlip must also rerun these analyses on the affected downstream section. FastFlip can reuse the results of these sub-analyses for all other sections. FastFlip uses the data-flow specification to identify such dependencies between inputs and outputs across sections.

Since the error injection analysis is the main contributor to analysis time, this approach significantly speeds up FastFlip compared to rerunning the error injection analysis on the full modified program, even when re-analyzing multiple sections.

4.8 Characteristics of Compatible Sub-Analyses

To enable its analysis, FastFlip must use error injection and SDC propagation analyses that satisfy certain key criteria:

Error injection analyses. The error injection analysis must separately report the outcome for errors in each error site in the program that the developer may wish to protect (e.g., [24, 67]), in contrast to just computing overall outcome statistics (e.g., [53]). This provides FastFlip with per-instruction error vulnerability information that is critical to its approach.

SDC propagation analyses. The SDC propagation analysis must support the SDC magnitude metric used by the error

injection and sensitivity analyses. The analysis must also support the propagation of SDCs whose magnitude is represented by a symbolic variable. Examples of such analyses are Chisel [47], DeepJ [34, 35], and Daisy [15].

Error and cost models. FastFlip’s formalism supports multiple error models. The error injection analysis may inject single or multi-bit errors into one or more error sites *within a single section*. The error sites can be individual instructions or coarser-grained program structures like statements. FastFlip also supports multiple cost models provided externally as a function $c(pc)$. This includes estimates of run time overhead for duplicating and comparing the results of single instructions (e.g., [56]), or the cost of specialized error detection for tasks or instruction blocks (e.g., [1, 2, 29]).

4.9 Factors That Affect the Precision of FastFlip

Inter-section masking. Inter-section masking occurs when an SDC present in one section is masked by a downstream section. FastFlip conservatively assumes that SDCs introduced in any section result in SDCs in the final outputs. The frequency of this masking is highly application-dependent.

Imprecision of sub-analyses. Imprecision in the error injection and SDC propagation sub-analyses used by FastFlip leads to imprecision in FastFlip. As FastFlip is a general approach that can use any sub-analysis that satisfies the requirements in Section 4.8, FastFlip’s precision can be improved by using newer, more precise sub-analyses as they become available.

Side effects. Due to errors, a section may cause additional, unexpected side effects that do not occur in error-free executions. Consequently, the section outputs may be SDC-free, but the error may still cause SDCs in later sections. For example, the error may cause the section to overwrite live data due to bad memory address calculations, or it may corrupt a live value while popping it from the stack at the end of the section. FastFlip mitigates these issues by checking all live variables at the end of each section for SDCs, and not just the section outputs.

Untested error sites. A small number of error sites in the program may not be included in any program section. For example, if sections are executed multiple times within an outer loop, then the instructions which check the loop exit conditions may be excluded from all program sections. FastFlip conservatively assumes that, if an error occurs at such an untested error site, then it will *always* produce an SDC-Bad outcome. More rigorously, FastFlip creates a special section s_{\perp} containing all such untested error sites j and assumes that $\forall j \in J_{s_{\perp}}, O_s(j) = (\infty, \dots, \infty)$. This reduces precision, as not all errors at the untested sites actually result in an SDC-Bad outcome.

4.10 Adapting and Compensating for Loss of Precision

A loss of precision due to the factors from Section 4.9 leads to a loss of *utility*. That is, it can cause FastFlip to protect against a smaller number of SDC-causing errors than expected, or increase the cost of protecting FastFlip’s selection of instructions beyond the minimum necessary cost. FastFlip adaptively

adjusts the target value v_{trgt} used in Section 4.6 to compensate for this loss of utility. In our experiments, this adjustment is insignificant except for one benchmark.

Measuring utility. FastFlip compares its utility to the utility obtained via a baseline monolithic error injection analysis that analyzes the whole program as a single section. FastFlip uses two primary metrics for measuring utility:

First, FastFlip treats the outcome labels of the baseline analysis as the ground truth and calculates the value of protecting its selection against SDC-Bad outcomes according to these alternate outcome labels. FastFlip refers to the protection value of its selection calculated in this manner as the *achieved value* v_{achv} . FastFlip then calculates the *loss of value* as $v_{loss} = v_{trgt} - v_{achv}$. v_{loss} measures the degree by which FastFlip undershoots v_{trgt} ; a lower v_{loss} is better.

Second, FastFlip calculates its excess cost over the baseline monolithic analysis. Specifically, if the costs associated with protecting the two selections of instructions against SDCs are c_{FF} (for FastFlip) and c_{Base} (for the baseline analysis), the excess cost is $c_{exc} = c_{FF} - c_{Base}$. c_{exc} measures the inefficiency of FastFlip’s selection for protecting against SDC-Bad outcomes compared to the baseline analysis’s selection; a lower c_{exc} is better.

When analyzing a program, FastFlip can simultaneously run the baseline error injection analysis for a minimal increase in analysis time. To do so, FastFlip checks the effect of each error in each section both on the section outputs and the final outputs. FastFlip efficiently calculates v_{loss} and c_{exc} using the outcome labels from FastFlip and the baseline analysis.

Adjusting the target value. FastFlip replaces the original target v_{trgt} with an adjusted target v'_{trgt} . Let the achieved value for this adjusted target be v'_{achv} . FastFlip minimizes v'_{trgt} such that $v'_{achv} \geq v_{trgt}$. If $v'_{trgt} > v_{trgt}$, then the cost of protecting FastFlip’s selection increases, with larger adjustments leading to larger increases. If instead $v'_{trgt} < v_{trgt}$, the cost decreases.

Target adjustment for modified program versions. If the number of modifications since the most recent target adjustment (m_{adj}) is below a developer-defined threshold (P_{adj}), FastFlip executes only its own time saving analysis and uses the existing adjusted target (v'_{trgt}) to choose the instructions to protect. As program modifications accumulate, the adjusted target may no longer provide the expected compensation for utility loss. Thus, once $m_{adj} \geq P_{adj}$, FastFlip recalculates v'_{trgt} by running a fresh analysis of the whole program while simultaneously running the monolithic analysis.

5 Methodology

5.1 Choice of Sub-Analyses

We instantiate FastFlip with the Approxilyzer [67] error injection analysis and the Chisel [47] SDC propagation analysis.

Approxilyzer is a bitflip error injection analysis that focuses on architectural CPU registers within each *dynamic* instruction in a program execution. Approxilyzer enumerates bitflip injection sites in the correct dynamic trace of

the program for a particular input. It uses heuristics to form equivalence classes of bitflips that cause similar outcomes. Then, it injects a bitflip for a single pilot from each equivalence class into the correct execution of the program within the gem5 simulator, continues the now tainted program execution (with possibly incorrect control flow), and records the effect of the bitflip on the program output. It then applies the outcome of this pilot bitflip to all members of the equivalence class.

Chisel is an SDC propagation analysis that calculates the end-to-end SDC propagation function $f_{T,\lambda}$ as a *conservative affine function* of the symbolic SDC variables $\varphi_{*,*}$. Chisel conservatively assumes that each program section always amplifies input SDCs by the maximum amplification factor for that section for any input. Chisel supports diverging control flow paths by calculating the maximum possible SDC amplification over any path. Due to these assumptions, it generates conservative end-to-end SDC specifications. We added support for symbolic SDC variables to Chisel in order to calculate symbolic end-to-end SDC specifications for FastFlip.

5.2 Error Model

While FastFlip supports multi-bit error models, our evaluation uses the same error model as Approxilyzer [67], described below, to ensure a fair comparison. We inject one single-bit transient error per simulation in an architectural general purpose or SSE2 register. We target both source and destination registers in dynamic instructions within the region of interest. We do not inject errors in special purpose, status, and control registers (e.g., %rsp, %rflags) as we assume that they always need protection which can be provided by hardware. Similarly, we assume that caches are protected by hardware (e.g., ECC). As in previous works (e.g., [37, 43]), we assume that the probability $p(j)$ that the error will occur at any error site j is uniform.

5.3 SDC Detection Model

We assume that an instruction selected for protection is duplicated and then followed by an equality check of the results. The duplicated code and increased register pressure leads to runtime overhead. However, by rearranging instructions and checks, the overhead/cost for *extensive* instruction duplication across the program can be reduced to 29% on average [48]. *Selective* duplication has even lower overhead (e.g. [32]).

Value and cost of detection. We adapt the value and cost model from [23]:

- The value $v(pc)$ of protecting a static instruction pc is proportional to the number of distinct errors injected in pc that produce an SDC-Bad outcome (using uniform $p(j)$).
- The cost $c(pc)$ of protecting pc is proportional to the number of dynamic instances of pc in the program trace.

5.4 Benchmarks

Table 1 presents our benchmarks, and we describe them next:

- *BScholes*: Black-Scholes analysis from PARSEC [5].

Table 1. List of FastFlip benchmarks. The Sections column shows *static*(\times *dynamic*) instances of sections in the trace.

Benchmark	Input size	Sections	# Error Sites (J)
BScholes	2 options	4 ($\times 2$)	36.7K
Campipe	32 \times 32	5 ($\times 1$)	72.7M
FFT	256 \times 2	5 ($\times 1$)	9.23M
LUD	16 \times 16	4 ($\times 2$)	1.75M
SHA2	32 bytes	3 ($\times 1$)	403K

- *Campipe*: The raw image processing pipeline for the Nikon D7000 camera from [71].
- *FFT*: Fast Fourier Transform from Splash-3 [57].
- *LUD*: Blocked LU decomposition from Splash-3 [57].
- *SHA2*: The SHA-256 hash function from [49].

For FFT and LUD, the input size is the same as the minimized input size found by Minotaur [45], a technique for reducing error injection time while retaining program counter coverage. For BScholes, we manually reduced the 21 option minimized input found by Minotaur down to 2 options without reducing program counter coverage. For Campipe, we use the reference 32 \times 32 input the implementation provided. For SHA2, we use a common cryptographic key size (256 bits).

5.5 Code Modifications for Benchmarks

To test the advantages offered by FastFlip for evolving programs, we also analyze modified versions of each benchmark. Then, we compare the results of the baseline analysis (must re-analyze the whole program) to those of FastFlip (must only inject errors in modified sections). We experiment with two types of semantics-preserving modifications:

Small modifications represent simple modifications that developers or compilers may make while optimizing and maintaining the program. Such modifications of up to 15 lines of code form a majority of open-source commits [3]. For Campipe and FFT, we store an expression used in multiple locations within the section into a variable to improve code readability. For LUD, we introduce a specialized version of a section that reduces the number of bounds checks if it detects that the matrix size is a multiple of the block size (as is the case for our input). For BScholes, we eliminate a redundant floating point operation in the cumulative normal distribution function. For SHA2, we similarly eliminate a redundant shift operation (without making the runtime input-dependent).

Large modifications replace a program section with a lookup table. The table maps inputs of that section to corresponding outputs. If the modified section finds the current input in this table, it returns the corresponding output. Otherwise, it executes the original section code.

5.6 Baseline, Comparison, and Experimental Setup

Software and hardware. FastFlip uses gem5-Approxilyzer version 22.1 [66] simulating an x86-64 CPU as the architecture simulator. We performed our experiments on AMD Epyc processors with 94 error injection experiment threads.

Region of interest. We focus on the computational portion of each benchmark and do not analyze I/O code.

SDC magnitude metric. We use the maximum element-wise absolute difference as the SDC metric. If $o_k[\ell]$ represents the ℓ^{th} element of an output o_k and the modified output due to an error is \hat{o}_k , then the SDC metric is $\max_{\ell} |o_k[\ell] - \hat{o}_k[\ell]|$.

SDC-Bad threshold. We first analyze all benchmarks assuming that any SDC is SDC-Bad ($\forall \lambda. \epsilon_{\lambda} = 0$). Next, we relax this requirement in Section 6.4 by assuming SDC magnitudes up to 0.01 to be tolerable, i.e., SDC-Good ($\forall \lambda. \epsilon_{\lambda} = 0.01$) for all benchmarks except SHA2 (whose applications require the output to be fully precise).

Sensitivity analysis parameters. As we consider the maximum tolerable SDC magnitude ϵ_{λ} to be 0.01 in Section 6.4, we use this as the maximum perturbation during the sensitivity analysis. To estimate the Lipschitz constant K (Equation 1), we perform 10^6 random perturbations up to ϵ_{λ} . For array inputs, we randomly perturb single, multiple, or all elements.

Comparison metrics. We compare the performance and utility of FastFlip to a baseline monolithic Approxilyzer-only approach. This baseline approach uses Approxilyzer to inject bitflips in the whole program at once and uses its results to select instructions to protect. For performance, we compare the analysis times of FastFlip and Approxilyzer run separately.

For comparing utility, we compare the selections of instructions to protect made by the two approaches using the value and cost metrics. For comparison, we choose three target values in the total value / cost trade-off space: $v_{\text{tgt}} \in \{0.90, 0.95, 0.99\}$, corresponding to protecting against 90%, 95%, and 99% of errors that cause unacceptably large SDCs.

Pruning error range. Approxilyzer’s use of equivalence classes as described in Section 5.1 speeds up both FastFlip and the baseline analysis. However, the pilot is not a perfect predictor of the outcomes for the pruned injections (i.e., the rest of the equivalence class). Figure 5 in Approxilyzer [67] shows that, on average, 4% of pruned injections have an outcome that significantly differs from that of the pilot.

Therefore, we establish an error range around the achieved value of SDC protection to account for this discrepancy among the outcomes of injections in an equivalence class. This error range depends on the pilot prediction inaccuracy and the fraction of error sites with SDC-Bad outcomes that are protected. For FFT, LUD, and BScholes, we use the benchmark-specific pilot prediction inaccuracy from Figure 5 in Approxilyzer [67] (3%, 4%, and 10% respectively). For Campipe and SHA2, we consider the average inaccuracy from the same figure (4%). We give details of the error range calculation in [27, Section 5.4.5]. If v_{achv} is within or above this error range around v_{tgt} , then we consider FastFlip’s result to be acceptable, even if $v_{\text{achv}} < v_{\text{tgt}}$.

Timeouts. FastFlip assumes that if the error causes the runtime of a program section to exceed $5 \times$ the nominal runtime, then the execution times out, which is a detected outcome. We use the same timeout rule for the Approxilyzer-only baseline.

6 Evaluation

6.1 Utility of FastFlip vs. Approxilyzer

Table 2 compares the utility of FastFlip and Approxilyzer for selective protection against SDCs, using the metrics described in Section 4.10. The pairs of columns show the utility comparison for the target protection values 0.90, 0.95, and 0.99 (90%, 95%, and 99% of SDC-causing errors) respectively. The first column in each pair shows FastFlip’s achieved protection value. The second column shows the cost of protecting FastFlip’s selection, and compares this to the cost of protecting Approxilyzer’s selection.

FastFlip successfully meets all target values for the unmodified (*None*) versions of each benchmark. Since FastFlip reuses the adjusted targets for the modified version it may not precisely meet the target for those modified versions. The maximum loss of value compared to the target is 0.017 (1.7%) for SHA2-Large. In all cases, the target value is within FastFlip’s value error range caused by injection pruning.

For most benchmarks, the cost of protecting FastFlip’s selection of instructions is at most 0.011 (1.1%) more than the cost of protecting Approxilyzer’s selection. The exception is Campipe, for which FastFlip’s cost is up to 0.068 (6.8%) higher. Unlike the other benchmarks, FastFlip has to aggressively adjust the target values for Campipe in order to meet the original targets to compensate for the loss of precision caused by inter-section masking. We observed that if we removed the last section of Campipe (the primary cause of inter-section masking), FastFlip’s target adjustments became less aggressive. This suggests that more precise SDC propagation analyses that also calculate the probability of SDC masking may reduce the need for target adjustment.

The geometric mean cost of protecting FastFlip’s selection is 0.601, 0.685, and 0.819 for the target protection values 0.90, 0.95, and 0.99, respectively. This shows that it is possible to protect against 90% of SDC-causing bitflips by protecting on average 60% of all dynamic instructions, but protecting against the remaining SDCs quickly leads to diminishing returns.

6.2 Performance of FastFlip vs. Approxilyzer

Table 3 compares the analysis time of FastFlip and Approxilyzer. Columns 1-2 show the benchmark name and version, respectively. Columns 3-4 show the analysis time for FastFlip and Approxilyzer *for that version of the benchmark*, respectively. Column 5 shows the speedup of FastFlip over Approxilyzer. We measure analysis time in *core-hours*. As the error injection analysis is highly parallelizable, the actual wall-clock time is much lower when using multiple CPU threads. The speedups in terms of wall-clock time have similar trends.

For FastFlip, error injection consumes 99% of the analysis time. The sensitivity analysis requires less than five minutes of wall-clock time. The symbolic SDC propagation analysis and knapsack solver each require under one minute, even for programs or inputs much larger than our benchmarks.

Table 2. Comparison of FastFlip and Approxilyzer utility when all SDCs are unacceptable (SDC-Bad) and target adjustment (Section 4.10) is used. A ✓ indicates that the achieved value is within the value error range of FastFlip.

Benchmark	Modification	$v_{\text{trgt}} = 0.90$		$v_{\text{trgt}} = 0.95$		$v_{\text{trgt}} = 0.99$	
		Value	Cost (diff)	Value	Cost (diff)	Value	Cost (diff)
BScholes	None	0.901✓	0.635 (+0.000)	0.950✓	0.717 (+0.000)	0.990✓	0.827 (+0.000)
	Small	0.899✓	0.634 (+0.003)	0.950✓	0.713 (+0.000)	0.990✓	0.821 (+0.000)
	Large	0.898✓	0.669 (+0.000)	0.949✓	0.753 (+0.000)	0.991✓	0.849 (+0.000)
Campipe	None	0.915✓	0.611 (+0.038)	0.950✓	0.676 (+0.017)	0.991✓	0.807 (+0.024)
	Small	0.924✓	0.611 (+0.060)	0.954✓	0.678 (+0.030)	0.990✓	0.807 (+0.034)
	Large	0.912✓	0.760 (+0.068)	0.961✓	0.819 (+0.043)	0.993✓	0.899 (+0.015)
FFT	None	0.900✓	0.544 (+0.011)	0.950✓	0.629 (+0.002)	0.990✓	0.780 (+0.000)
	Small	0.904✓	0.542 (+0.010)	0.950✓	0.629 (+0.004)	0.990✓	0.781 (+0.002)
	Large	0.900✓	0.492 (+0.001)	0.950✓	0.586 (−0.000)	0.987✓	0.716 (−0.016)
LUD	None	0.900✓	0.603 (+0.000)	0.950✓	0.694 (+0.000)	0.990✓	0.873 (+0.000)
	Small	0.901✓	0.606 (+0.002)	0.951✓	0.698 (+0.002)	0.990✓	0.875 (+0.001)
	Large	0.902✓	0.560 (+0.002)	0.951✓	0.640 (+0.003)	0.990✓	0.826 (−0.001)
SHA2	None	0.900✓	0.666 (+0.001)	0.950✓	0.772 (+0.000)	0.990✓	0.908 (+0.001)
	Small	0.900✓	0.665 (+0.000)	0.949✓	0.771 (−0.001)	0.990✓	0.908 (+0.000)
	Large	0.883✓	0.476 (−0.007)	0.943✓	0.551 (−0.003)	0.985✓	0.655 (−0.007)

Table 3. Analysis execution time comparison.

Bench.	Modif.	Analysis time (core-hours)		
		FastFlip	Approxilyzer	Speedup
BScholes	None	69 hrs	65 hrs	0.9×
	Small	42 hrs	62 hrs	1.5×
	Large	3 hrs	24 hrs	8.4×
Campipe	None	2459 hrs	2631 hrs	1.1×
	Small	158 hrs	2720 hrs	17.2×
	Large	45 hrs	494 hrs	11.0×
FFT	None	980 hrs	520 hrs	0.5×
	Small	300 hrs	509 hrs	1.7×
	Large	93 hrs	513 hrs	5.5×
LUD	None	694 hrs	602 hrs	0.9×
	Small	80 hrs	625 hrs	7.8×
	Large	94 hrs	441 hrs	4.7×
SHA2	None	726 hrs	728 hrs	1.00×
	Small	718 hrs	726 hrs	1.01×
	Large	43 hrs	45 hrs	1.05×

To enable target adjustment, FastFlip simultaneously runs the Approxilyzer analysis as described in Section 4.10. We use the methodology from [45, Section 4.7] to confirm that the time required for this approach is at most 1% more than the greater of the analysis times of FastFlip and Approxilyzer for the unmodified versions of the benchmarks. As FastFlip reuses the adjusted targets for modified benchmarks, it does not need to use this approach when the benchmarks are modified.

The two approaches have similar analysis times for the unmodified (*None*) versions of all benchmarks except FFT. For

FFT, Approxilyzer prunes a larger number of injections since it finds that an operation is repeated in different program sections. As FastFlip injects errors into each section independently, it cannot similarly prune injections across sections.

For the modified benchmarks, the speedup of FastFlip depends on the number of error sites that FastFlip must re-analyze compared to the full program. If the modified program sections represent a small fraction of the total error sites, then FastFlip provides large speedups. Critically, FastFlip is at least 1.7× faster when analyzing the modified versions of FFT. If the modified program sections represent a large fraction of the total error sites, then FastFlip provides smaller speedups. This leads to the negligible speedups for SHA2, where we modified the most expensive section of the program.

These results show that FastFlip can save significant time when analyzing evolving programs. Here, even a single re-analysis helped to offset the original analysis overhead. For modern software systems that developers gradually modify over time, FastFlip provides ever increasing savings.

6.3 Effects of Target Value Adjustment

For all benchmarks except Campipe, the original and adjusted target values are virtually the same: the difference is within 0.4% of the original targets. As such, the conclusions presented in Section 6.1 are valid even without target adjustment for these benchmarks. For Campipe, target adjustment helps to address the issue with the last section described in Section 6.1.

Table 4 compares the utility of FastFlip and Approxilyzer for Campipe when FastFlip does not use target adjustment. The format is similar to that of Table 2, except that we omit the cost columns and focus on whether FastFlip still achieves the

Table 4. Comparison of FastFlip and Approxilyzer utility for Campipe *without* target adjustment. A ✓ indicates that the achieved value is within the value error range of FastFlip, while a ✗ indicates the opposite. Table 2 shows the improved results *with* target adjustment.

Benchmark	Modif.	Value @ $v_{\text{trgt}} =$		
		0.90	0.95	0.99
Campipe	None	0.848✗	0.920✓	0.977✓
	Small	0.879✓	0.925✓	0.980✓
	Large	0.868✗	0.925✗	0.979✓

target value. Without target adjustment, FastFlip undershoots the targets by as much as 0.052 (5.2%) and the original target values do not always fall within FastFlip’s achieved value error range. These results show the importance of target adjustment for ensuring that FastFlip meets the original protection target.

6.4 Ignoring Acceptably Small SDCs

We next compare the utility of FastFlip and Approxilyzer when small SDCs (≤ 0.01) are considered acceptable (SDC-Good) and the analyses focus on protecting against errors that cause larger SDCs (SDC-Bad).

FastFlip successfully meets all target values for all benchmarks. The maximum loss of value compared to the target is 0.014 (1.4%) for FFT-Large. In all cases, the target value is within FastFlip’s achieved value error range caused by injection pruning. For most benchmarks, the cost of protecting FastFlip’s selection of instructions is at most 0.020 (2%) more than the cost of protecting Approxilyzer’s selection. The exception is Campipe, for which FastFlip’s cost is higher by as much as 0.057 (5.7%), again due to aggressive target adjustment.

The geomean cost of protecting FastFlip’s selection is 0.619, 0.720, and 0.849 for the target protection values 0.90, 0.95, and 0.99, respectively. FastFlip obtains the results for this scenario at the same time as the results in Table 2 for negligible additional analysis time (< 1 minute). We describe these results in more detail in [27, Section 5.5.5].

7 Limitations

In Section 4.6, FastFlip assumes that the cost of protecting multiple instructions is equal to the sum of protecting each individual instruction in that set. However, protecting multiple adjacent instructions via techniques such as instruction duplication may lead to excess basic block fragmentation or register pressure, which increases protection cost beyond the sum of the protection cost for each instruction in isolation. DRIFT [48] describes methods for mitigating this issue.

Section 4.8 describes the criteria that error injection and SDC propagation analyses must satisfy for use with FastFlip. As these analyses form a core part of FastFlip’s approach, they affect FastFlip’s precision and the error models that it can support. Section 4.9 describes several factors that reduce the

precision of FastFlip, as well as methods to mitigate these issues. Not all of these issues can be completely eliminated, reducing FastFlip’s precision compared to the baseline monolithic analysis. Our evaluation shows that FastFlip effectively compensates for this loss of precision via target adjustment. However, target adjustment can lead to an increase in protection cost, such as for Campipe in Table 2. As more powerful, precise, and general error injection and SDC propagation analyses become available, FastFlip will be able to use them to support more error models and provide more precise results.

Section 4.10 describes a simple heuristic that FastFlip uses to determine if it needs to re-analyze the whole program after a modification for more precise target adjustment. Our evaluation shows that this heuristic is generally capable of maintaining FastFlip’s precision for modified programs. Regardless, we believe that more complex heuristics (e.g., those based on the lines of code modified) could provide better precision by accounting for the size and nature of the modification.

Section 6.2 shows that FastFlip is much slower than Approxilyzer when analyzing the unmodified version of FFT due to less effective injection pruning. We expect that FastFlip will also exhibit such a slowdown for other computations in which injection pruning is particularly effective. However, Table 3 also shows that this initial disadvantage is quickly amortized when analyzing modified versions of such programs.

8 Related Work

Error injection analyses. Error injection analyses operate at different levels of abstraction, including hardware, assembly, and IR [11, 12, 18, 28, 30, 31, 39–41, 43, 53, 54, 59]. These analyses typically use *sampling*: they select a statistically significant number of error sites at random and only perform error injections at those sites. While this is sufficient for providing overall outcome statistics, a developer cannot use such results to determine which specific instructions or blocks of instructions are particularly vulnerable to SDCs in order to protect them. However, FastFlip can still use these analyses if they are modified to perform full instruction-level error injection like Approxilyzer [66, 67]. Minotaur [45] reduces the size of inputs (and therefore analysis time) required to test the reliability of programs when subjected to errors without compromising on the coverage of error sites. FastFlip further reduces analysis times for these minimized inputs as the program evolves.

Li et al. [40] show that error injection at higher levels of abstraction does not easily model the impact of all lower level hardware errors. Similarly, Papadimitriou and Gizopoulos [52] show that injecting errors in various SRAM hardware structures gives different results compared to injecting errors at higher levels of abstraction. AVGI [53] builds on [52] to show that hardware errors manifest in software in different ways, but result in similar distributions of final outcomes across applications. Santos et al. [58] similarly examine how faults injected at the RTL level affect common GPU instructions, and inject these effects into applications at the software level

to provide overall outcome statistics and identify vulnerable hardware components. Unfortunately, such analysis techniques that aim to efficiently determine the effect of low-level faults via hybrid fault injection are too slow even for small program sizes when the outcomes are needed for *each* error site for fine-grained software-based SDC protection. If techniques such as AVGI become scalable for analyzing each error site in the future, we believe FastFlip would be able to use them.

Reliability analyses without error injection. ePVF [19] is a dynamic analysis which finds locations where a bitflip will cause a crash, as opposed to an SDC, with $\sim 90\%$ accuracy. TRIDENT [38] uses empirical observations of error propagation in programs to predict the overall SDC probability of a program and the SDC probabilities of individual instructions. Several other works [8, 44, 60] use analytical modeling to detect SDCs in a program. While these analyses can be faster than error injection analyses, they are less accurate and may not be able to precisely estimate the magnitude of the output SDC due to an error. FastFlip’s compositional nature makes error injection analysis more affordable by amortizing the cost of analyzing evolving programs over time.

SDC propagation analyses. SDC propagation analyses either propagate SDCs forward through programs [7, 9, 15, 21, 35], or propagate SDC bounds backwards through programs [22, 47]. While we used the Chisel [47] SDC propagation analysis to evaluate FastFlip, it can use other analyses that satisfy the conditions described in Section 4.8.

Mutlu et al. [50] predict the effect of bitflips injected into iterative applications on the final output by analyzing the effects of fault injections on a limited number of iterations. While it may give an advantage over FastFlip for applications that iterate the same operation multiple times, unlike FastFlip, it cannot handle programs with multiple sections that perform distinct operations, such as our benchmarks.

Hardware-based selective protection. Researchers have examined the use of selective hardware hardening (e.g., via redundancy or ECC) for improving hardware reliability while limiting the use of additional chip area [14, 42, 55, 73]. These techniques find and replicate only those hardware components that, as a result of transient errors, produce unacceptable outcomes across the range of typical applications for that hardware. FastFlip efficiently provides information which can be used to apply *additional*, software-based selective protection tailored for specific applications, as opposed to adding further hardware protections irrelevant to other applications.

Software-based selective SDC protection. Unlike crashes, timeouts, or clearly invalid data, SDCs are more difficult to detect by nature. SWIFT [56] uses instruction duplication to detect errors in computational instructions. To reduce overhead, it makes use of downtime in a program’s instruction schedule. DRIFT [48] further reduces overhead by clustering together checks of multiple duplicated instructions to reduce basic block fragmentation. SWIFT and DRIFT aim to *completely* eliminate the possibility of SDCs occurring due to single

bitflip errors in the duplicated computational instructions. nZDC [17] provides comparable overhead to SWIFT while also protecting programs from 99.6% of SDCs caused by single bitflip errors during load, store, and control flow instructions.

Shoestring [20] finds and duplicates only particularly vulnerable instructions. Hari et al. [23] propose protecting blocks of instructions with single detectors placed at the end of loops or function calls. These two techniques use the results of error injection analyses to guide selective instruction duplication. Coarse-grained approaches place detectors at the task level [1, 2, 25, 29, 72]. We consider such techniques to be *clients* of FastFlip. They can provide FastFlip with information about the runtime overhead of protecting various instructions or instruction blocks. In return, FastFlip can provide precise information on which instructions should be protected in order to minimize runtime overhead while protecting against a developer-defined fraction of SDC-causing errors. After these techniques protect FastFlip’s selection of instructions, FastFlip can re-analyze the protected sections to confirm the decrease in SDC vulnerability. For FastFlip, we focused on efficiently handling program modifications in general. Testing code modifications specifically designed to reduce SDC vulnerability is an interesting topic for future work.

Incremental program analysis. Incremental techniques have a rich history in improving run-time of common program analyses that study control-flow equivalence and/or complex heap data structure properties, e.g., [26, 33, 36, 51, 62] or ML model robustness [64, 65]. In contrast, FastFlip incrementally analyzes the impact of hardware errors on the program execution, which are out of scope of off-the-shelf incremental techniques that operate on coarser-grained program properties.

9 Conclusion

We presented FastFlip, the first systematic approach that combines error injection and SDC propagation analyses to enable fast error injection analysis of evolving programs. When developers modify the program, FastFlip’s compositional nature allows it to selectively re-analyze only the modified code sections, making it 3.2 \times faster on average (geomean) and up to 17.2 \times faster compared to a baseline non-compositional analysis that re-analyzes the whole program. Additionally, the value and cost of protecting FastFlip’s selection of instructions closely tracks that of the baseline analysis.

FastFlip can reduce the burden of repeated error injection analysis whenever developers fix program bugs, add optimizations, and add protections for vulnerable instructions. FastFlip therefore represents the first step toward including resiliency analysis and hardening as first-class citizens in the standard software development workflow, which continually compiles and tests software after each code modification.

Acknowledgments

The research in this paper was supported in part by NSF Grants CCF-1846354, CCF-1956374, and CCF-2217144.

References

- [1] Sara Achour and Martin C. Rinard. 2015. Approximate computation with outlier detection in Topaz. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications* (Pittsburgh, PA, USA) (*OOPSLA 2015*). 711–730. <https://doi.org/10.1145/2814270.2814314>
- [2] P. Agrawal. 1988. Fault tolerance in multiprocessor systems without dedicated redundancy. *IEEE Trans. Comput.* 37, 3 (1988), 358–362. <https://doi.org/10.1109/12.2174>
- [3] Abdulkareem Alali, Huzefa Kagdi, and Jonathan I. Maletic. 2008. What's a Typical Commit? A Characterization of Open Source Software Repositories. In *2008 16th IEEE International Conference on Program Comprehension*. 182–191. <https://doi.org/10.1109/ICPC.2008.24>
- [4] Rizwan A. Ashraf, Roberto Gioiosa, Gokcen Kestor, Ronald F. DeMara, Chen-Yong Cher, and Pradip Bose. 2015. Understanding the propagation of transient errors in HPC applications. In *SC '15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–12. <https://doi.org/10.1145/2807591.2807670>
- [5] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. 2008. The PARSEC benchmark suite: characterization and architectural implications. In *Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques* (Toronto, Ontario, Canada) (*PACT '08*). 72–81. <https://doi.org/10.1145/1454115.1454128>
- [6] S. Borkar. 2005. Designing reliable systems from unreliable components: the challenges of transistor variability and degradation. *IEEE Micro* 25, 6 (2005), 10–16. <https://doi.org/10.1109/MM.2005.110>
- [7] James Bornholt, Todd Mytkowicz, and Kathryn S. McKinley. 2014. Uncertain<T>: a first-order type for uncertain data. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems* (Salt Lake City, Utah, USA) (*ASPLOS '14*). 51–66. <https://doi.org/10.1145/2541940.2541958>
- [8] Brett Boston, Zoe Gong, and Michael Carbin. 2018. Leto: verifying application-specific hardware fault tolerance with programmable execution models. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 163 (Oct. 2018), 30 pages. <https://doi.org/10.1145/3276533>
- [9] Brett Boston, Adrian Sampson, Dan Grossman, and Luis Ceze. 2015. Probability type inference for flexible approximate programming. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications* (Pittsburgh, PA, USA) (*OOPSLA 2015*). 470–487. <https://doi.org/10.1145/2814270.2814301>
- [10] Dan G. Cacuci and Mihaela Ionescu-Bujor. 2004. A Comparative Review of Sensitivity and Uncertainty Analysis of Large-Scale Systems—II: Statistical Methods. *Nuclear Science and Engineering* 147, 3 (2004), 204–217. <https://doi.org/10.13182/04-54CR>
- [11] Jon Calhoun, Luke Olson, and Marc Snir. 2014. FlipIt: An LLVM Based Fault Injector for HPC. In *Revised Selected Papers, Part I, of the Euro-Par 2014 International Workshops on Parallel Processing - Volume 8805*. 547–558. https://doi.org/10.1007/978-3-319-14325-5_47
- [12] Chun-Kai Chang, Sangkug Lym, Nicholas Kelly, Michael B. Sullivan, and Mattan Erez. 2018. Hamartia: A Fast and Accurate Error Injection Framework. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*. 101–108. <https://doi.org/10.1109/DSN-W.2018.00046>
- [13] Swarat Chaudhuri, Sumit Gulwani, Roberto Lubliner, and Sara Navidpour. 2011. Proving programs robust. In *Proceedings of the 19th ACM SIGSOFT Symposium and the 13th European Conference on Foundations of Software Engineering* (Szeged, Hungary) (*ESEC/FSE '11*). 102–112. <https://doi.org/10.1145/2025113.2025131>
- [14] Josie E. Rodriguez Condia, Paolo Rech, Fernando Fernandes dos Santos, Luigi Carrot, and Matteo Sonza Reorda. 2021. Protecting GPU's Microarchitectural Vulnerabilities via Effective Selective Hardening. In *2021 IEEE 27th International Symposium on On-Line Testing and Robust System Design (IOLTS)*. 1–7. <https://doi.org/10.1109/IOLTS52814.2021.9486703>
- [15] Eva Darulova, Anastasiia Izycheva, Fariha Nasir, Fabian Ritter, Heiko Becker, and Robert Bastian. 2018. Daisy - Framework for Analysis and Optimization of Numerical Programs (Tool Paper). In *Tools and Algorithms for the Construction and Analysis of Systems*. 270–287.
- [16] Eva Darulova and Viktor Kuncak. 2014. Sound compilation of reals. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (San Diego, California, USA) (*POPL '14*). 235–248. <https://doi.org/10.1145/2535838.2535874>
- [17] Moslem Didehban and Aviral Shrivastava. 2016. nZDC: A compiler technique for near Zero Silent Data Corruption. In *2016 53rd ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6. <https://doi.org/10.1145/2897937.2898054>
- [18] Waleed Dweik, Murali Annavam, and Michel Dubois. 2014. Reliability-Aware Exceptions: Tolerating intermittent faults in microprocessor array structures. In *2014 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 1–6. <https://doi.org/10.7873/DATE.2014.114>
- [19] Bo Fang, Qining Lu, Karthik Pattabiraman, Matei Ripeanu, and Sudhanva Gurumurthi. 2016. ePVF: An Enhanced Program Vulnerability Factor Methodology for Cross-Layer Resilience Analysis. In *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 168–179. <https://doi.org/10.1109/DSN.2016.24>
- [20] Shuguang Feng, Shantanu Gupta, Amin Ansari, and Scott Mahlke. 2010. Shoestring: probabilistic soft error reliability on the cheap. In *Proceedings of the Fifteenth International Conference on Architectural Support for Programming Languages and Operating Systems* (Pittsburgh, Pennsylvania, USA) (*ASPLOS XV*). 385–396. <https://doi.org/10.1145/1736020.1736063>
- [21] Vimuth Fernando, Keyur Joshi, Jacob Laurel, and Sasa Misailovic. 2023. Diamont: dynamic monitoring of uncertainty for distributed asynchronous programs. *Int. J. Softw. Tools Technol. Transf.* 25, 4 (Nov. 2023), 521–539. <https://doi.org/10.1007/s10009-023-00717-y>
- [22] Vimuth Fernando, Keyur Joshi, and Sasa Misailovic. 2019. Verifying safety and accuracy of approximate parallel programs via canonical sequentialization. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 119 (Oct. 2019), 29 pages. <https://doi.org/10.1145/3360545>
- [23] Siva Kumar Sastry Hari, Sarita V. Adve, and Helia Naeimi. 2012. Low-cost program-level detectors for reducing silent data corruptions. In *IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2012)*. 1–12. <https://doi.org/10.1109/DSN.2012.6263960>
- [24] Siva Kumar Sastry Hari, Sarita V. Adve, Helia Naeimi, and Pradeep Ramachandran. 2012. Relyzer: exploiting application-level fault equivalence to analyze application resiliency to transient faults. In *Proceedings of the Seventeenth International Conference on Architectural Support for Programming Languages and Operating Systems* (London, England, UK) (*ASPLOS XVII*). 123–134. <https://doi.org/10.1145/2150976.2150990>
- [25] Hukerikar and Engelmann. 2017. Resilience Design Patterns: A Structured Approach to Resilience at Extreme Scale. *Supercomput. Front. Innov.: Int. J.* 4, 3 (Sept. 2017), 4–42. <https://doi.org/10.14529/jsfi170301>
- [26] Kenneth Johnson, Radu Calinescu, and Shinji Kikuchi. 2013. An incremental verification framework for component-based software systems. In *Proceedings of the 16th International ACM Sigsoft Symposium on Component-Based Software Engineering* (Vancouver, British Columbia, Canada) (*CBSE '13*). 33–42. <https://doi.org/10.1145/2465449.2465456>
- [27] Keyur Joshi. 2024. *Compositional Analysis of the Effects of Uncertainty on Computations*. PhD dissertation. University of Illinois, Urbana-Champaign. Available at <https://hdl.handle.net/2142/124160>.
- [28] Keyur Joshi, Vimuth Fernando, and Sasa Misailovic. 2019. Statistical Algorithmic Profiling for Randomized Approximate Programs. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 608–618. <https://doi.org/10.1109/ICSE.2019.00071>
- [29] Keyur Joshi, Vimuth Fernando, and Sasa Misailovic. 2020. Aloe: verifying reliability of approximate programs in the presence of recovery mechanisms. In *Proceedings of the 18th ACM/IEEE International Symposium on Code Generation and Optimization* (San Diego, CA, USA)

- (CGO '20), 56–67. <https://doi.org/10.1145/3368826.3377924>
- [30] Manolis Kaliorakis, Dimitris Gizopoulos, Ramon Canal, and Antonio Gonzalez. 2017. MeRLiN: Exploiting dynamic instruction behavior for fast and accurate microarchitecture level reliability assessment. In *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*. 241–254. <https://doi.org/10.1145/3079856.3080225>
- [31] Manolis Kaliorakis, Sotiris Tselonis, Athanasios Chatzidimitriou, Nikos Foutiris, and Dimitris Gizopoulos. 2015. Differential Fault Injection on Microarchitectural Simulators. In *2015 IEEE International Symposium on Workload Characterization*. 172–182. <https://doi.org/10.1109/ISWC.2015.28>
- [32] Daya Shanker Khudia and Scott Mahlke. 2014. Harnessing Soft Computations for Low-Budget Fault Tolerance. In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*. 319–330. <https://doi.org/10.1109/MICRO.2014.33>
- [33] Y. Lakhnech, S. Bensalem, S. Berezin, and S. Owre. 2001. Incremental Verification by Abstraction. In *Tools and Algorithms for the Construction and Analysis of Systems*. 98–112.
- [34] Jacob Laurel, Siyuan Brant Qian, Gagandeep Singh, and Sasa Misailovic. 2023. Synthesizing Precise Static Analyzers for Automatic Differentiation. *Proc. ACM Program. Lang.* 7, OOPSLA2, Article 291 (Oct. 2023), 29 pages. <https://doi.org/10.1145/3622867>
- [35] Jacob Laurel, Rem Yang, Gagandeep Singh, and Sasa Misailovic. 2022. A dual number abstraction for static analysis of Clarke Jacobians. *Proc. ACM Program. Lang.* 6, POPL, Article 56 (Jan. 2022), 30 pages. <https://doi.org/10.1145/3498718>
- [36] Steven Lauterburg, Ahmed Sobeih, Darko Marinov, and Mahesh Viswanathan. 2008. Incremental state-space exploration for programs with dynamically allocated data. In *2008 ACM/IEEE 30th International Conference on Software Engineering*. 291–300. <https://doi.org/10.1145/1368088.1368128>
- [37] Guanpeng Li and Karthik Pattabiraman. 2018. Modeling Input-Dependent Error Propagation in Programs. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 279–290. <https://doi.org/10.1109/DSN.2018.00038>
- [38] Guanpeng Li, Karthik Pattabiraman, Siva Kumar Sastry Hari, Michael Sullivan, and Timothy Tsai. 2018. Modeling Soft-Error Propagation in Programs. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 27–38. <https://doi.org/10.1109/DSN.2018.00016>
- [39] Jianli Li and Qingping Tan. 2013. SmartInjector: Exploiting intelligent fault injection for SDC rate analysis. In *2013 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFTS)*. 236–242. <https://doi.org/10.1109/DFT.2013.6653612>
- [40] Man-Lap Li, Pradeep Ramachandran, Ulya R. Karpuzcu, Siva Kumar Sastry Hari, and Sarita V. Adve. 2009. Accurate microarchitecture-level fault modeling for studying hardware faults. In *2009 IEEE 15th International Symposium on High Performance Computer Architecture*. 105–116. <https://doi.org/10.1109/HPCA.2009.4798242>
- [41] Xiaodong Li, Sarita V. Adve, Pradip Bose, and Jude A. Rivers. 2008. Online Estimation of Architectural Vulnerability Factor for Soft Errors. In *2008 International Symposium on Computer Architecture*. 341–352. <https://doi.org/10.1109/ISCA.2008.9>
- [42] F. Libano, B. Wilson, J. Anderson, M. J. Wirthlin, C. Cazzaniga, C. Frost, and P. Rech. 2019. Selective Hardening for Neural Networks in FPGAs. *IEEE Transactions on Nuclear Science* 66, 1 (2019), 216–222. <https://doi.org/10.1109/TNS.2018.2884460>
- [43] Qining Lu, Mostafa Farahani, Jiesheng Wei, Anna Thomas, and Karthik Pattabiraman. 2015. LLFI: An Intermediate Code-Level Fault Injection Tool for Hardware Faults. In *2015 IEEE International Conference on Software Quality, Reliability and Security*. 11–16. <https://doi.org/10.1109/QRS.2015.13>
- [44] Qining Lu, Guanpeng Li, Karthik Pattabiraman, Meeta S. Gupta, and Jude A. Rivers. 2017. Configurable Detection of SDC-causing Errors in Programs. *ACM Trans. Embed. Comput. Syst.* 16, 3, Article 88 (March 2017), 25 pages. <https://doi.org/10.1145/3014586>
- [45] Abdulrahman Mahmoud, Radha Venkatagiri, Khaliq Ahmed, Sasa Misailovic, Darko Marinov, Christopher W. Fletcher, and Sarita V. Adve. 2019. Minotaur: Adapting Software Testing Techniques for Hardware Errors. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (Providence, RI, USA) (ASPLOS '19)*. 1087–1103. <https://doi.org/10.1145/3297858.3304050>
- [46] Kaisa Miettinen. 1998. *A Priori Methods*. 115–129. https://doi.org/10.1007/978-1-4615-5563-6_5
- [47] Sasa Misailovic, Michael Carbin, Sara Achour, Zichao Qi, and Martin C. Rinard. 2014. Chisel: reliability- and accuracy-aware optimization of approximate computational kernels. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications (Portland, Oregon, USA) (OOPSLA '14)*. 309–328. <https://doi.org/10.1145/2660193.2660231>
- [48] Konstantina Mitropoulou, Vasileios Porpodas, and Marcelo Cintra. 2014. DRIFT: Decoupled CompileR-Based Instruction-Level Fault-Tolerance. In *Languages and Compilers for Parallel Computing*. 217–233.
- [49] Alain Mosnier. 2023. SHA-2 algorithm implementations. <https://github.com/amosnier/sha-2>.
- [50] Burcu O. Mutlu, Gokcen Kestor, Adrian Cristal, Osman Unsal, and Sriram Krishnamoorthy. 2019. Ground-Truth Prediction to Accelerate Soft-Error Impact Analysis for Iterative Methods. In *2019 IEEE 26th International Conference on High Performance Computing, Data, and Analytics (HiPC)*. 333–344. <https://doi.org/10.1109/HiPC.2019.00048>
- [51] Peter W. O'Hearn. 2018. Continuous Reasoning: Scaling the impact of formal methods. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (Oxford, United Kingdom) (LICS '18)*. 13–25. <https://doi.org/10.1145/3209108.3209109>
- [52] George Papadimitriou and Dimitris Gizopoulos. 2021. Demystifying the System Vulnerability Stack: Transient Fault Effects Across the Layers. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. 902–915. <https://doi.org/10.1109/ISCA52012.2021.00075>
- [53] George Papadimitriou and Dimitris Gizopoulos. 2023. AVGI: Microarchitecture-Driven, Fast and Accurate Vulnerability Assessment. In *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. 935–948. <https://doi.org/10.1109/HPCA56546.2023.10071105>
- [54] Konstantinos Parasyris, Georgios Tziantzoulis, Christos D. Antonopoulos, and Nikolaos Bellas. 2014. GemFI: A Fault Injection Tool for Studying the Behavior of Applications on Unreliable Substrates. In *2014 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. 622–629. <https://doi.org/10.1109/DSN.2014.96>
- [55] Ilija Polian and John P. Hayes. 2011. Selective Hardening: Toward Cost-Effective Error Tolerance. *IEEE Design & Test of Computers* 28, 3 (2011), 54–63. <https://doi.org/10.1109/MDT.2010.120>
- [56] G.A. Reis, J. Chang, N. Vachharajani, R. Rangan, and D.I. August. 2005. SWIFT: software implemented fault tolerance. In *International Symposium on Code Generation and Optimization*. 243–254. <https://doi.org/10.1109/CGO.2005.34>
- [57] Christos Sakalis, Carl Leonardsson, Stefanos Kaxiras, and Alberto Ros. 2016. Splash-3: A properly synchronized benchmark suite for contemporary research. In *2016 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 101–111. <https://doi.org/10.1109/ISPASS.2016.7482078>
- [58] Fernando F. dos Santos, Josie E. Rodriguez Condia, Luigi Carro, Matteo Sonza Reorda, and Paolo Rech. 2021. Revealing GPUs Vulnerabilities by Combining Register-Transfer and Software-Level Fault Injection. In *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 292–304. <https://doi.org/10.1109/DSN48987.2021.00042>

- [59] Horst Schirmeier, Martin Hoffmann, Christian Dietrich, Michael Lenz, Daniel Lohmann, and Olaf Spinczyk. 2015. FAIL²: An Open and Versatile Fault-Injection Framework for the Assessment of Software-Implemented Hardware Fault Tolerance. In *2015 11th European Dependable Computing Conference (EDCC)*. 245–255. <https://doi.org/10.1109/EDCC.2015.28>
- [60] Vilas Sridharan and David R. Kaeli. 2009. Eliminating microarchitectural dependency from Architectural Vulnerability. In *2009 IEEE 15th International Symposium on High Performance Computer Architecture*. 117–128. <https://doi.org/10.1109/HPCA.2009.4798243>
- [61] Phillip Stanley-Marbell, Armin Alaghi, Michael Carbin, Eva Darulova, Lara Dolecek, Andreas Gerstlauer, Ghayoor Gillani, Djordje Jevdjic, Thierry Moreau, Mattia Cacciotti, Alexandros Daglis, Natalie Enright Jerger, Babak Falsafi, Sasa Misailovic, Adrian Sampson, and Damien Zufferey. 2020. Exploiting Errors for Efficiency: A Survey from Circuits to Applications. *ACM Comput. Surv.* 53, 3, Article 51 (June 2020), 39 pages. <https://doi.org/10.1145/3394898>
- [62] Benno Stein, Bor-Yuh Evan Chang, and Manu Sridharan. 2021. Demanded abstract interpretation. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (Virtual, Canada) (PLDI 2021)*. 282–295. <https://doi.org/10.1145/3453483.3454044>
- [63] Anna Thomas and Karthik Pattabiraman. 2013. Error detector placement for soft computation. In *2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 1–12. <https://doi.org/10.1109/DSN.2013.6575353>
- [64] Shubham Ugare, Debangshu Banerjee, Sasa Misailovic, and Gagandeep Singh. 2023. Incremental Verification of Neural Networks. *Proc. ACM Program. Lang.* 7, PLDI, Article 185 (June 2023), 26 pages. <https://doi.org/10.1145/3591299>
- [65] Shubham Ugare, Gagandeep Singh, and Sasa Misailovic. 2022. Proof transfer for fast certification of multiple approximate neural networks. *Proc. ACM Program. Lang.* 6, OOPSLA1, Article 75 (April 2022), 29 pages. <https://doi.org/10.1145/3527319>
- [66] Radha Venkatagiri, Khaliq Ahmed, Abdulrahman Mahmoud, Sasa Misailovic, Darko Marinov, Christopher W. Fletcher, and Sarita V. Adve. 2019. gem5-Approxilyzer: An Open-Source Tool for Application-Level Soft Error Analysis. In *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 214–221. <https://doi.org/10.1109/DSN.2019.00033>
- [67] Radha Venkatagiri, Abdulrahman Mahmoud, Siva Kumar Sastry Hari, and Sarita V. Adve. 2016. Approxilyzer: Towards a systematic framework for instruction-level approximate computing and its application to hardware resiliency. In *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 1–14. <https://doi.org/10.1109/MICRO.2016.7783745>
- [68] Shaobu Wang, Guangyan Zhang, Junyu Wei, Yang Wang, Jiesheng Wu, and Qingchao Luo. 2023. Understanding Silent Data Corruptions in a Large Production CPU Population. In *Proceedings of the 29th Symposium on Operating Systems Principles (Koblenz, Germany) (SOSP '23)*. 216–230. <https://doi.org/10.1145/3600006.3613149>
- [69] Wikipedia. 2024. Continuous Integration. https://en.wikipedia.org/wiki/Continuous_integration.
- [70] Graham R Wood and BP Zhang. 1996. Estimation of the Lipschitz constant of a function. *Journal of Global Optimization* 8 (1996), 91–103.
- [71] Yuan Yao. 2023. CAVA: Camera Vision Pipeline on gem5-Aladdin. <https://github.com/yaoyuannnn/cava>.
- [72] A. Ziv and J. Bruck. 1997. Performance optimization of checkpointing schemes with task duplication. *IEEE Trans. Comput.* 46, 12 (1997), 1381–1386. <https://doi.org/10.1109/12.641939>
- [73] Christian G. Zoellin, Hans-Joachim Wunderlich, Ilia Polian, and Bernd Becker. 2008. Selective Hardening in Early Design Steps. In *2008 13th European Test Symposium*. 185–190. <https://doi.org/10.1109/ETS.2008.30>