

# Fine Grained Dataflow Tracking with Proximal Gradients

Gabriel Ryan<sup>†</sup>, Abhishek Shah<sup>†</sup>, Dongdong She<sup>†</sup>, Koustubha Bhat<sup>‡</sup>, Suman Jana<sup>†</sup>  
<sup>†</sup>*Columbia University*, <sup>‡</sup>*Vrije Universiteit Amsterdam*

## Abstract

Dataflow tracking with Dynamic Taint Analysis (DTA) is an important method in systems security with many applications, including exploit analysis, guided fuzzing, and side-channel information leak detection. However, DTA is fundamentally limited by the Boolean nature of taint labels, which provide no information about the significance of detected dataflows and lead to false positives/negatives on complex real world programs.

We introduce proximal gradient analysis (PGA), a novel, theoretically grounded approach that can track more accurate and fine-grained dataflow information. PGA uses proximal gradients, a generalization of gradients for non-differentiable functions, to precisely compose gradients over non-differentiable operations in programs. Composing gradients over programs eliminates many of the dataflow propagation errors that occur in DTA and provides richer information about how each measured dataflow effects a program.

We compare our prototype PGA implementation to three state of the art DTA implementations on 7 real-world programs. Our results show that PGA can improve the F1 accuracy of data flow tracking by up to 33% over taint tracking (20% on average) without introducing any significant overhead (< 5% on average). We further demonstrate the effectiveness of PGA by discovering 22 bugs (20 confirmed by developers) and 2 side-channel leaks, and identifying exploitable dataflows in 19 existing CVEs in the tested programs.

## 1 Introduction

Dataflow analysis with dynamic taint analysis (DTA) is a fundamental building block in many common systems security tasks, such as automated vulnerability analysis, guided fuzzing, discovering information leaks, and malware analysis [4, 14, 20, 38, 42, 56]. DTA analyzes dataflow between a specified set of sources and sinks in a program

by instrumenting the program and tracking taint as it executes [30, 34].

However, DTA is fundamentally limited by the Boolean information contained in taint labels: data either is tainted by a given source or not; there are no intermediate states or other sources of information. This means there is no way to identify and prioritize which dataflows are most significant. For example, given a series of operations  $x1 = a * 8$ ;  $x2 = b / 8$ ;  $y = x1 + x2$ ; changes to the value of  $a$  will have a larger effect on the value of  $y$  than changes to value of  $b$ , but taint labels cannot make this distinction. Moreover, it limits the ability of DTA frameworks to account for dataflows that are dependent on how operations compose. For example, in  $x1 = x * 2$ ;  $x2 = x1 \& 1$ ; variable  $x2$  will only be affected by changes in the first bit of  $x1$ , but changes to  $x$  will not affect  $x2$  due to the intermediate multiplication by 2.

While most DTA systems incorporate some special rules to handle these types of cases, we find in our evaluation (Section 5.2.1) that current DTA systems with these rules still make many errors in predicting dataflows, even at high compiler optimization levels that eliminate most intermediate operations. These errors have prevented DTA from being successfully applied in applications such as detecting keyloggers and memory corruption attacks [6, 12, 48, 49].

The limitations of DTA led several researchers to propose Quantitative Information Flow (QIF) based methods as a more fine grained form of dataflow [33]. However, while QIF is able to track data more precisely, computing these measures is computationally expensive and does not scale effectively to large programs [26].

In this paper, we propose an alternate measure of dataflow that addresses the limitations of DTA while retaining its advantages in scalability. We observe that **gradient**, a multi-variate generalization of derivatives from elementary calculus, is a popular method for tracking the influence of inputs through differentiable models [16]. In particular, gradients have been used in neural

```

taint source: x1, x2, x3
taint sink: y

// input: x1, x2, x3 = 10
int y = x1 + x2*x2 - 100*x3;
if (y > THRESH) {
    // vulnerability
}

```

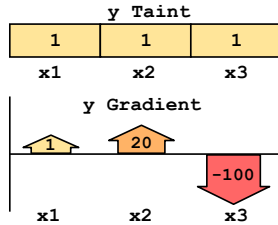


Figure 1: Example program in which gradient can guide a search to reach a vulnerability. While taint tracking identifies  $y$  as tainted by all three inputs, gradient measures the magnitude and direction of each influence, identifying that  $x_3$  is the most influential input and that minimizing it will maximize  $y$  due to its negative gradient.

networks to perform a variety of tasks that are analogous to the applications of DTA in program analysis, including generating inputs to trigger errors, explaining output behaviors, and maximizing test coverage [5, 21, 37, 46, 47, 51].

The additional information provided by gradients confer two crucial advantages: (i) **Fine-grained tracking**. Gradients measure both the *magnitude* and *direction* of influence, which indicate how changes to an operation’s input will effect its output. This means gradients can be used to identify which marked sources are most influential, and how they will effect program behavior. This is illustrated in Figure 1, in which the magnitude of the gradient identifies the most influential input, and the direction of the gradient indicates how that input can be changed to reach a vulnerability. (ii) **Precise composition**. Gradients can be used to identify when an operation input will have no effect on its output due to composition. For  $x_1 = x * 2$ ;  $x_2 = x_1 \& 1$ ; the gradient of  $x_1$  will be 2 and the gradient of  $x_2$  will be 0, which correctly identifies that the first bit will never change in the operation  $x_1 \& 1$  and therefore there will be no dataflow.

However, in general, programs contain many non-differentiable operations with different types of non-smooth behavior (e.g. bitwise operations, integer arithmetic, and branches as shown in Figure 2) that cannot be differentiated directly. Therefore, we build on the rich non-smooth calculus literature to define generalized gradients for programs that satisfy weaker forms of chain rule [22, 31, 52]. To evaluate generalized gradients on programs, we use **proximal gradients**, which compute gradient on non-differentiable operations by finding the local minima [36]. Proximal gradients provide a theoretically grounded framework for gradient evaluation that allows us to precisely track dataflow across real-world programs with minimal compositional errors.

We implement a prototype of Proximal Gradient Analysis (PGA) as an LLVM pass that instruments programs during compilation to compute proximal gradients. We

compare PGA to three state-of-the-art DTA systems on 7 widely used applications and show that PGA achieves up to 33% better F1 accuracy (20% on average) than DataFlowSanitizer, the best performing DTA system, without incurring any significant (<5%) extra overhead. We apply PGA to guided fuzzing and show that using PGA achieves up to 56% higher edge coverage (10% on average) than DTA in a controlled comparison, as well as improving the coverage achieved by a state-of-the-art fuzzer NEUZZ by 13% on average [45]. Finally, we use PGA to discover 22 bugs and 2 side-channel leaks, and analyze 19 existing CVEs.

The rest of this paper is organized as follows. First, Section 2 summarizes the background on different generalizations of gradients to non-smooth analysis. Next, we describe our methodology for computing proximal gradients on real-world programs in Section 3. We describe the details of our implementation of proximal gradient analysis in Section 4, Section 5 contains the details of our evaluation setup and results, and we discuss the tradeoffs of PGA and DTA in Section 6. Finally, we summarize related work in Section 7 and conclude in Section 8.

Our main contributions are:

1. We are the first, to the best of our knowledge, to use non-smooth analysis for dataflow tracking in real-world programs. Specifically, we design, implement, and evaluate Proximal Gradient Analysis (PGA), a novel, theoretically grounded technique for measuring fine grained influence in real-world programs.
2. We implement our PGA framework for automatically computing and tracking proximal gradients as an LLVM pass. An open source release of PGA is available at <https://github.com/gryan11/PGA>.
3. We perform extensive experimental evaluation of PGA and compare it to three state of the art DTA implementations, DataFlowSanitizer, libdft, and Neutaint, on 7 popular, real-world programs. PGA achieves up to 33% higher F1 accuracy than DTA (20% on average) without introducing significant additional overhead (on average <5%). PGA also achieves up to 56% improvement in new edge coverage relative to DTA (10% on average) for data-flow-guided fuzzing, as well as improving the coverage achieved by a state-of-the-art fuzzer NEUZZ by 12.9% on average.
4. We demonstrate that PGA’s fine-grained tracking is helpful for finding and analyzing different types of bugs and information leaks. In our experiments, PGA found 22 bugs and 2 side-channel leaks in our tested programs. PGA also detected the exploitable dataflow in 19 known CVEs, including 2 where DTA fails.

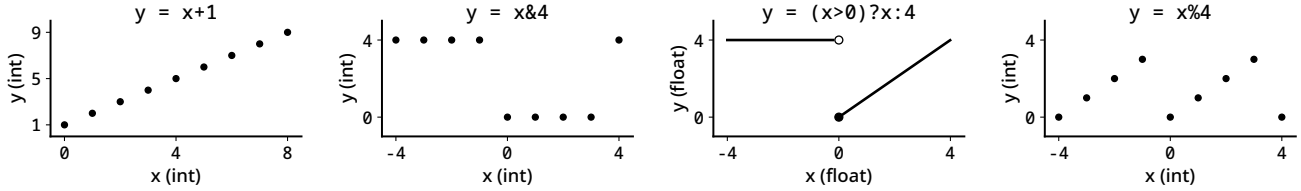


Figure 2: Different types of discrete and discontinuous operations that occur in real-world programs

## 2 Background

Our approach to gradient-based dataflow analysis draws on several techniques from the mathematical analysis and optimization literature. We provide a summary of the relevant methods below. We first summarize standard methods for computing gradients over compositions of smooth functions, and then review techniques from the non-smooth analysis literature that can be applied to computing gradients over programs.

### 2.1 Smooth Analysis

**Gradients.** The derivative for a smooth scalar function  $f(x)$  is defined as  $f'(x) = \lim_{\delta x \rightarrow 0} \frac{f(x+\delta x) - f(x)}{\delta x}$ , where  $f: \mathbb{R} \rightarrow \mathbb{R}$ . If a function has a derivative for all points in its domain, then it is considered a differentiable function. The gradient is a generalization of the derivative to multivariate functions, where  $f: \mathbb{R}^n \rightarrow \mathbb{R}$  and  $\nabla f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ , that can be understood as the slope of the function at the point where it is evaluated. When a function is vector-valued (i.e.  $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ ), the Jacobian generalizes gradient by evaluating the gradient of each of the  $m$  outputs:  $\mathcal{J}f: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ . For the rest of the paper, functions are multi-variate unless otherwise noted.

**Chain Rule.** Gradients of compositions of differentiable functions can be computed from gradients of the individual functions. This is known as the chain rule of calculus and is defined as follows, where  $\circ$  indicates the composition of two functions  $f$  and  $g$ , and  $f'$  and  $g'$  are their respective gradients:

$$(f \circ g)' = (f' \circ g) * g' \quad (1)$$

Elementwise multiplication is used when  $f$  and  $g$  are multivariate.

**Automatic Differentiation.** Automatic Differentiation (AutoDiff) uses the chain rule to compute the gradient for potentially large compositions of differentiable functions. AutoDiff has been a longstanding tool in computational modeling and is a core component of deep learning frameworks such as Tensorflow [2, 53]. However, existing AutoDiff methods and frameworks are limited to working with mostly continuous functions with limited discontinuity (e.g. ReLUs in neural networks).

### 2.2 Non-smooth Analysis

Extensive work has been done in the field of mathematical analysis on methods for approximating gradients over non-smooth functions. In this section we consider general multivariate functions of type  $f: \mathbb{R}^n \rightarrow \mathbb{R}$ . We first describe a generalized type of continuity, called Lipschitz continuity, that applies to non-smooth operations in programs, and then define a generalization of gradients that apply to Lipschitz continuous functions.

**Lipschitz Continuity.** A function is Lipschitz continuous if its output does not change too much for small changes in the input. Formally, a function  $f$  is Lipschitz continuous if there exists a constant  $K$  (called the Lipschitz constant) that bounds how much the value of  $f$  can change between any two points in its domain. Figure 3a shows a simple Lipschitz continuous function along with the corresponding Lipschitz constant. In general the operations in any useful computation will yield a Lipschitz continuous function.

**Generalized Gradients.** On Lipschitz continuous functions, *generalized gradients* are used to approximate gradients [13, 40]. Generalized gradients consist of generalized directional derivatives, which evaluate the gradient in a single direction as shown in Figure 3b. A generalized directional derivative in a direction  $\mathbf{v} \in \mathbb{R}^n$  is defined as follows:

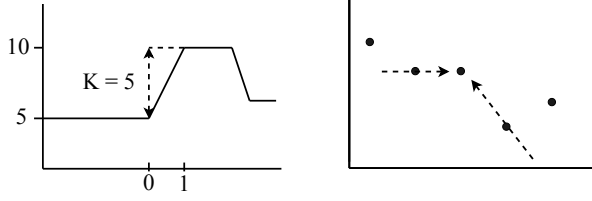
$$f'(x; \mathbf{v}) = \lim_{y \rightarrow x, \lambda \downarrow 0} \sup \frac{f(y + \lambda \mathbf{v}) - f(y)}{\lambda} \quad (2)$$

Here  $x$  and  $y$  are two points in the domain of  $f$  where  $x$  is the point the derivative is evaluated, and  $\lambda$  is a distance along the vector  $\mathbf{v}$  that the derivative is taken in. The chain rule for directional derivatives with functions  $g: \mathbb{R}^n \rightarrow \mathbb{R}^n$  and  $f: \mathbb{R}^n \rightarrow \mathbb{R}$  is defined:

$$(f \circ g)'(x; \mathbf{v}) = f'(g(x); g'(x; \mathbf{v})) \quad (3)$$

When applied to generalized directional derivatives the composing functions must be monotonic. Several relaxed versions of the chain rule apply to generalized derivatives under different weaker assumptions about the composite functions [22, 31, 52].

A generalized gradient is approximated with a set of directional derivatives based on a matrix  $\mathbf{V} \in \mathbb{R}^{n \times p} =$



(a) Lipshitz func. (b) Directional deriv.

Figure 3: **Example of a Lipschitz function with  $K=5$  and directional derivatives on a discrete function.**

$[\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_p]$  of  $p$  vectors in the domain of  $f$  representing the directions in which the derivatives are evaluated.

$$f'(x; \mathbf{V}) = [f'(x; \mathbf{v}_1), f'(x; \mathbf{v}_2), \dots, f'(x; \mathbf{v}_p)] \quad (4)$$

When  $f$  is a composition of functions, the chain rule from Eq. 3 can be applied to each of the generalized directional derivatives:

$$(f \circ g)'(x; \mathbf{V}) = [(f \circ g)'(x; \mathbf{v}_1), \dots, (f \circ g)'(x; \mathbf{v}_p)] \quad (5)$$

### 3 Methodology

At a high level, our gradient propagation framework, PGA, is similar to Autodiff, computing the gradient of each operation and using the results as inputs to the next gradient computation. However, unlike Autodiff, we approximate the gradients of discrete functions with proximal gradients.

**Proximal Gradients.** Since programs are generally composed of discrete operations on integers, we define a gradient approximation called proximal gradients that can be evaluated on these discrete functions. Proximal gradients use the minima of a function within a nearby region defined with a special operator called the proximal operator [36]. This can be evaluated on both discrete and continuous functions  $f: X^n \rightarrow X$ , where  $X$  is a set with euclidean norm that can represent integers or floats.

$$\text{prox}_f(x) = \arg \min_y (f(y) + \frac{1}{2} \|x - y\|_2^2) \quad (6)$$

The notation  $\arg \min_y$  indicates that the operator evaluates to the value of  $y$  that minimizes the sum of the function  $f(y)$  and the distance cost.

We use the proximal operator to compute each generalized directional derivative  $f'(x; \mathbf{v})$ . Given a function  $f$  representing a program operation, we constrain the proximal operator from Eq. 6 to a direction  $\mathbf{v}$ :

$$\text{prox}_f(x; \mathbf{v}) = \arg \min_y (f(y) + \frac{1}{2} \|x - y\|_2^2) \quad (7)$$

$$\text{where } y = x + t\mathbf{v} : t \in \mathbb{N}, y \in X^n$$

We then define the proximal directional derivative based on the difference with  $\text{prox}_f(x; \mathbf{v})$  constrained in the direction  $\mathbf{v}$  and scaled by the direction magnitude  $\|\mathbf{v}\|_2$ :

$$\text{prox}'_f(x; \mathbf{v}) = \frac{f(\text{prox}_f(x; \mathbf{v})) - f(x)}{\|\text{prox}_f(x; \mathbf{v}) - x\|_2} * \|\mathbf{v}\|_2 \quad (8)$$

This takes the same form as the generalized directional derivative (Eq. 2), but evaluated with the proximal operator. A proximal gradient is defined for a set of direction vectors  $V$  like the generalized gradient (Eq. 4) using proximal directional derivatives:

$$\text{prox}'_f(x; \mathbf{V}) = [\text{prox}'_f(x; \mathbf{v}_1), \dots, \text{prox}'_f(x; \mathbf{v}_p)] \quad (9)$$

Using proximal gradients allows us to evaluate gradients on discrete operations in programs as if they were continuous nonsmooth functions and apply the associated chain rule for generalized gradients in Eq. 5. For the rest of this paper, we refer to ‘proximal gradients’ simply as ‘gradients’ unless otherwise specified.

### 3.1 Program Gradient Evaluation

To compute gradients over programs with PGA, we model a program as a discrete function  $P: X^n \rightarrow X^n$ , and model the program state  $x \in X^n$  as a vector (e.g.  $x$  could model a byte array of size  $n$  representing the program memory and registers).  $P$  is composed of  $N$  functions  $P_i: X^n \rightarrow X^n, i \in \{1..N\}$  representing individual operations on the program state:

$$P(x) = P_N \circ P_{N-1} \circ \dots \circ P_2 \circ P_1(x)$$

Each program operation  $P_i$  is modeled as a combination of  $n$  non-smooth scalar valued functions  $f_{ij}: X^n \rightarrow X, j \in \{1..n\}$  that define how  $P_i$  modifies each variable in the program state.

$$P_i(x) = [f_{i,1}(x), f_{i,2}(x), \dots, f_{i,n}(x)]$$

We evaluate each  $f'_{ij}$  in  $P'_i$  using the proximal directional derivative (Eq. 8):

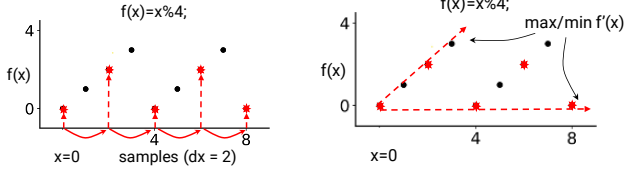
$$P'_i(x; \mathbf{v}) = [\text{prox}'_{f_{i,1}}(x; \mathbf{v}), \dots, \text{prox}'_{f_{i,n}}(x; \mathbf{v})]$$

To compose derivatives for a given operation  $P_i$  from the previous operation  $P_{i-1}$ , we individually compose the derivatives of each  $f_{ij}$  in  $P_i$  from the previous operation  $P_{i-1}$ :

$$(P_i \circ P_{i-1})'(x; \mathbf{v}) = [(f_{i,1} \circ P_{i-1})'(x; \mathbf{v}), \dots, (f_{i,n} \circ P_{i-1})'(x; \mathbf{v})] \quad (10)$$

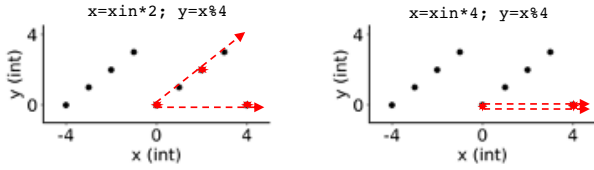
where each  $(f_{ij} \circ P_{i-1})'(x; \mathbf{v})$  is defined based on the directional derivative chain rule in Eq. 5:

$$(f_{ij} \circ P_{i-1})'(x; \mathbf{v}) = f'_{ij}(P_{i-1}(x); P'_{i-1}(x; \mathbf{v}))$$



(a) Sample with derivative  $dx/dinput=2$ . (b) Compute min/max  $f'(x)$  from samples.

Figure 4: **Derivative sampling procedure on an  $x\%4$  operation where the  $x$  derivative wrt. input  $dx/dinput=2$ . Samples are first collected at intervals of 2 and then used to compute the max/min directional derivative.**



(a) Composition of mul 2 with mod 4. When  $dx/dxin=2$ ,  $dy/dx=1$ . (b) Composition of mul 4 with mod 4. When  $dx/dxin=4$ ,  $dy/dx=0$ .

Figure 5: **Proximal Derivative evaluation on composition of a mul and mod operation at  $xin=0$ , with samples in red. The step size for the proximal derivative on  $x\%4$  is determined by the derivative  $dx/dxin$ . In subfigure (b), when  $xin$  is first multiplied by 4,  $dx/dxin=4$  and the sample step size for  $x\%4$  is 4. This causes the proximal derivative to evaluate to 0, which correctly indicates there is no dataflow over  $x\%4$  after mul by 4.**

Using the chain rule from Eq. 10, we can compute a directional derivative for each final state of the program  $P$  by chaining derivatives of the individual operations.

$$P'(x; \mathbf{v}) = (P_N \circ P_{N-1} \circ \dots \circ P_2 \circ P_1)'(x; \mathbf{v})$$

We then compute the proximal gradient using Eq. 9 for each program state by combining derivatives for a set of direction vectors represented by a matrix  $\mathbf{V}$ :

$$P'(x; \mathbf{V}) = [P'(x; \mathbf{v}_1), P'(x; \mathbf{v}_2), \dots, P'(x; \mathbf{v}_p)]$$

This is the same approach used in Automatic Differentiation, but extended to discrete functions and generalized gradients. This chained gradient approximation is designed to be error-free for all locally Lipschitz convex functions as well as some locally Lipschitz non-convex functions that meet the requirements for the non-smooth chain rule (e.g., monotonicity).

### 3.2 Proximal Derivative Evaluation

When applying proximal directional derivatives in practice, we make two modifications to the proximal direc-

#### Algorithm 1 Proximal Derivative computation on a non-smooth operation.

<b>Input:</b>	$op$	$\leftarrow$ program operation
	$x1, x2$	$\leftarrow$ operation inputs
	$dx1, dx2$	$\leftarrow$ $x1, x2$ components of $\mathbf{v}$
	$N$	$\leftarrow$ maximum samples

```

1: if  $dx1 = 0$  and  $dx2 = 0$  then
2:   return  $dy \leftarrow 0$ 
3: end if
4:  $y \leftarrow op(x1, x2)$ 
5: initialize size  $N$  arrays  $S$  and  $S_{cost}$ 
6: for  $i = 1$  to  $N$  do
7:    $x1_i \leftarrow x1 + dx1 * i$ 
8:    $x2_i \leftarrow x2 + dx2 * i$ 
9:    $y_i \leftarrow op(x1_i, x2_i)$ 
10:   $distance_i^2 \leftarrow (x1 - x1_i)^2 + (x2 - x2_i)^2$ 
11:  add  $-|y - y_i| + \frac{1}{2} distance_i^2$  to  $S_{cost}$  array
12:  add  $y_i$  to  $S$  array
13: end for
14:  $i_{prox} \leftarrow$  index of min sample in  $S_{cost}$ 
15:  $y_{prox} \leftarrow$  recover sample  $i_{prox}$  from  $S$ 
16: return  $dy \leftarrow (y_{prox} - y) / i_{prox}$ 

```

tional derivative defined in Eq. 8 to model program behavior more closely.

First, we only consider the inputs to the operation itself in a function  $f^* : X^k \rightarrow X, k \in \{1..n\}$  and associated  $v^*$ , where  $k$  is the number of inputs to the operation. To simplify notation we drop the  $*$ , and for the rest of the paper assume  $f$  and  $\mathbf{v}$  to refer to their  $k$  dimensional variants on the current operation.

Second, we modify the proximal operator to select a nearby point that maximizes *absolute change* in  $f$ , which we denote  $|\delta f|$ :

$$\text{prox}_{|\delta f|}(x; \mathbf{v}) = \underset{y}{\text{argmin}} \left( -|f(x) - f(y)| + \frac{1}{2} \|x - y\|_2^2 \right) \quad (11)$$

$$\text{where } y = x + t * \mathbf{v} : t \in \mathbb{N}, y \in X^n$$

This modified proximal operator selects the largest generalized derivative of  $f$  based on either the maximum or minimum of  $f$  in the direction  $\mathbf{v}$  (these correspond to the *supremum* or *infimum* of a generalized derivative). Accounting for both is necessary in dataflow analysis to avoid missing possible dataflows.

**Proximal Derivative Algorithm.** Algorithm 1 defines how we compute the proximal derivative for an operation  $op$  that has two input variables  $x1$  and  $x2$ , and returns an output  $y$ . We denote the derivatives of the inputs  $x1$  and  $x2$  and output  $y$  to be  $dx1$ ,  $dx2$ , and  $dy$ , where  $dx1$  and  $dx2$  are components of  $\mathbf{v}$ , and  $dy$  is computed using the proximal derivative with a maximum



sample budget  $N$ . The same algorithm can be applied to operations with any number of inputs from 1 to  $n$  by adjusting the number input variables. Figure 4 shows an example of the proximal derivative procedure being applied to a  $x\%4$  operation.

We observed that when the proximal gradient is nonzero, it almost always uses a point within a few samples of the current point due to rapid increase of the proximal cost term  $distance^2$  in the proximal operator. Therefore, we set  $N$  to a small constant (5 in our evaluation), and evaluate the proximal derivative in that range.

Figure 5 gives an example of evaluating Algorithm 1 on a non-smooth operation  $y = x\%4$ . When the input is multiplied by 2 as in Figure 5a, the algorithm samples at intervals of 2 and evaluates a derivative of 1 based on the maximum absolute difference ( $|\delta f|$ ) measure. However, when the input is multiplied by 4 as in Figure 5b, the algorithm samples at intervals of 4 and evaluates a derivative of 0 because the samples are all 0. This 0 derivative indicates that the composition of functions  $x=x\text{in}4$ ;  $y=x\%4$  will always have the same output and therefore has no dataflow.

### 3.3 Derivative Propagation Rules

We define a general framework for propagating derivatives over 5 abstract classes of operations that need to be handled in program analysis: floating point operations, integer valued operations, loading and storing variables, branching, and function calls to external libraries.

1. **Floating point operations:** We treat floating point operations as continuous functions and apply the standard chain rule (Eq. 10) with their analytic derivatives. If there are any potentially non-smooth floating point operations, such as floating point modulo, or typecasting between floating point types, we use proximal derivatives.
2. **Integer operations:** We consider any boolean or typecasting involving integers to be integer operations, as well as any arithmetic, bit shifting, or modulo on integer or pointer types. In general we use proximal derivatives on all integer operations, although in some cases such as arithmetic addition and multiplication we use analytic derivatives as an optional optimization.
3. **Load and Store:** When variables are stored or loaded from memory, their associated derivatives are also stored or loaded (our implementation uses shadow memory to track derivatives in memory, although any associative tracking mechanism could be used). If the memory address passed to a load

instruction has a nonzero derivative, we set the derivative of the loaded variable to 1.0 if it does not already have a nonzero derivative. This is a simplifying approximation that may lead errors in evaluating the proximal gradient. However, we note that proximal derivatives on load operations can potentially be evaluated by sampling adjacent memory locations. We leave this to future work.

4. **Branches:** When dynamically computing derivatives, we can only reason about the derivative on the current execution path. If computing a derivative would require sampling an alternate execution path, we instead set that derivative to 0. Therefore, when a branch is encountered, we set any derivatives to 0 that are based on samples that would change the branch condition. This approach may miss some parts of the gradient but ensures we do not propagate incorrect derivatives. We note that sampling across multiple execution paths when handling branches could yield more accurate proximal derivatives and reason about control flow data flows (i.e. implicit data flows), we leave this to future work.
5. **External Library Functions:** Provided they do not have side effects, derivatives on external library function calls can be computed using proximal derivatives, while functions with side effects must be handled on a case by case basis. When an external function overwrites a buffer, we also clear the stored derivatives associated with that buffer.

### 3.4 Program Gradient as Dataflow

To use gradients as a measure of dataflow, we compute gradient between a set of user defined sources and sinks. We set the initial vectors in  $\mathbf{V}$  so that each vector is all 0s except for an initial derivative on each source of +1 or -1. We then execute the program and propagate the derivatives over each operation with the chain rule and derivatives defined in Algorithm 1. While the program is executing we record derivatives at each sink, and accumulate the gradient on each sink from all the sources. Cumulatively, the gradients on all sinks form the Jacobian  $\mathcal{J}$  between sources and sinks.

Algorithm 2 formally describes the process for computing the gradients from a set of sources to each designated sink in program. The returned Jacobian  $\mathcal{J}$  contains the gradients of each sink based on the largest derivative propagated to it from each source (sinks may record multiple derivatives from a single source if, for example, the sink is in a loop).

---

**Algorithm 2 Program Gradient Evaluation.**


---

<b>Input:</b>	$P$ $\leftarrow$ program under analysis
	$x$ $\leftarrow$ program input
	$Sources$ $\leftarrow$ $n$ dataflow sources
	$Sinks$ $\leftarrow$ $m$ dataflow sinks

```

1: initialize  $\mathbf{V}$  to empty set  $\{\}$ 
2: initialize  $\mathcal{J}$   $n \times m$  matrix to 0s
3: for  $src$  in  $Sources$  do
4:    $\mathbf{v}_{src}^+ \leftarrow [dsrc = 1, \text{otherwise } 0]$ 
5:    $\mathbf{v}_{src}^- \leftarrow [dsrc = -1, \text{otherwise } 0]$ 
6:   add  $\mathbf{v}_{src}^+$  and  $\mathbf{v}_{src}^-$  to  $\mathbf{V}$ 
7: end for
8: Execute  $P$  on input  $x$ , tracking  $P'(x; \mathbf{V})$ 
9: for  $sink$  in  $Sinks$  do
10:  for each recorded  $\frac{dsink}{dsrc}$  do
11:   if  $|\frac{dsink}{dsrc}| > |\mathcal{J}[src, sink]|$  then
12:     $\mathcal{J}[src, sink] \leftarrow \frac{dsink}{dsrc}$ 
13:   end if
14:  end for
15: end for
16: return  $\mathcal{J}$ 

```

---

## 4 Implementation

We implement PGA as a new sanitizer in the LLVM framework [27] called Gradient Sanitizer (**grsan**). We use LLVM because it allows us to instrument a program during compilation after it has been converted to LLVM’s intermediate representation. This means that **grsan** can be used to instrument any program written in a language supported by LLVM, and incurs lower runtime overhead than binary instrumentation frameworks such as PIN or Valgrind [1, 32]. However, we note that PGA could also be implemented in a binary instrumentation framework to facilitate analysis in cases where source code is not available.

**Overall Architecture.** We base **grsan** on LLVM’s taint tracking implementation, DataFlowSanitizer (**dfsan**), which uses shadow memory to track taint labels. For each byte of application memory, there are two corresponding bytes of shadow memory that store the taint label for that byte.

We modify **dfsan** in the following two ways: First, we add additional metadata associated with each label that stores the gradient information, which is stored in a separate table as shown in Figure 6. Each label in the shadow memory is associated with a distinct derivative value in the gradient table. The 0 label is reserved for 0 derivative, and any shadow memory lookup on a constant or unlabeled variable returns label 0.

Second, we change the dataflow propagation rules to compute gradients over each operation. Figure 6 shows an

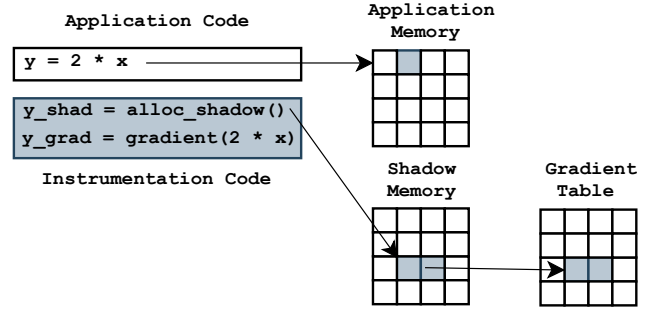


Figure 6: **grsan** architecture illustrating how proximal gradients are propagated.

example of how the **grsan** instrumentation works. Given an operation  $y=2*x$ , the instrumentation first looks up the derivative for each input, 2 and  $x$ , from shadow memory. If any input has a nonzero derivative, it computes the derivative for the output  $y$  and generates a new shadow memory label by incrementing the current max label by 1. It then allocates space in the shadow memory and gradient table and stores the new label and associated derivative of  $y$ .

As an additional optimization, when storing an operation’s output derivative we first compare it to the input derivatives. If the output derivative is equal to either, we apply the label of the equivalent input derivative to the output instead of generating a new label and gradient table entry. Since many operations do not change the value of the derivative (e.g.  $x = x+1$ ), this significantly reduces the number of distinct labels that need to be tracked.

In the current implementation, **grsan** tracks derivatives from a single source at a time, propagating the two derivatives from the source in parallel. When computing a gradient over multiple sources (e.g. bytes in an input file), we execute the program once for each source. We intend to extend **grsan** to support multiple sources in parallel in future work.

**Gradient Propagation Instrumentation.** For differentiable operations such as a floating point multiplication (**fmul**), **grsan** uses the analytical derivative of the operation. For nondifferentiable operations such as bitwise **And**, **grsan** uses an optimized version of proximal derivatives from Algorithm 1 that returns the first nonzero derivative it encounters when sampling. We found this approximation picked the same values that the proximal operator would select and is computationally lighter (i.e. does not require computing exponents).

We leave most external function calls uninstrumented, but some operations in **glibc** are given special instrumentation. We set the gradients for any buffer overwritten by **fread** or **memset** to 0, and the gradients of buffers copied by **memcpy** or **strcpy** are also copied. Type cast-

Library	Test Command	SLOC	File Format
zlib-1.2.11	<code>minigzip -d</code>	3228	GZ/ZIP
libjpeg-9c	<code>djpeg</code>	8,857	JPEG
mupdf-1.14.0	<code>mutool show</code>	123,562	PDF
libxml2-2.9.7	<code>xmllint</code>	73,920	XML
binutils-2.30	<code>objdump -xD</code>	72,955	ELF
	<code>strip</code>	56,330	
	<code>size</code>	52,991	ELF

Table 1: **Test programs used in our evaluation.**

ing instructions are handled by simply copying labels from the original value to the result.

## 5 Evaluation

We evaluate PGA by comparing its performance directly to DTA, and in direct applications for bug finding and security analysis. Specifically, we run experiments to answer the following questions:

1. **Dataflow Accuracy:** Is PGA more accurate than DTA in tracking dataflows?
2. **Overhead:** How does the overhead introduced by PGA compare to DTA?
3. **Guided Fuzzing:** Does using PGA to guide fuzzing lead to better edge coverage?
4. **CVE Analysis:** Can PGA detect and analyze recent CVEs that taint is typically used to detect?
5. **Bug Discovery:** Is PGA an effective tool for finding bugs?
6. **Information Leaks:** Can PGA detect and analyze memory and timing-based information leaks?

### 5.1 Experimental Setup

**Test Programs.** We perform tests on a set of 5 widely used file parsing libraries and 7 total programs. We use file parsers because these programs often must process files from untrusted sources, making them a common target for attacks. Table 1 shows the test programs and SLOC associated with each executable tested. In total the programs have 391,883 SLOC.

**Fuzzers Evaluated.** For our fuzzing experiments, we use the latest version of NEUZZ <sup>1</sup> and VUZZER <sup>2</sup>.

**Test Environment.** All of our evaluations are performed on an Ubuntu 16.04 server with an Intel Xeon E5-2623 v4 2.60GHz CPU and 192G of memory unless otherwise specified.

<sup>1</sup>[www.github.com/Dongdongshe/neuzz](https://www.github.com/Dongdongshe/neuzz)

<sup>2</sup>[www.github.com/vusec/vuzzer64](https://www.github.com/vusec/vuzzer64)

### 5.2 Performance

We first evaluate the performance of PGA as a tool for dynamic dataflow analysis. In our experiments, we compare PGA to DataFlowSanitizer (`dfsan`), LLVM’s state-of-the-art DTA implementation. Since our implementation of PGA is based on the `dfsan` architecture, our setup ensures that any differences in performance between PGA and DTA are due to the respective performance of gradient and taint and not due to differences in the underlying architectures.

We compare performance in three areas: first, we estimate the accuracy of the dataflows predicted by PGA and DTA. Second, we evaluate the overhead introduced by the PGA instrumentation. Third, we compare the edge coverage achieved by a dataflow-guided fuzzer using either PGA or DTA to guide its mutation strategy.

**Evaluation Inputs.** We use the same set of initial input files for all of the performance evaluations. The `gzip`, `pdf`, and `ELF` files are sourced from the AFL sample seeds included in the distribution<sup>3</sup>. The `jpeg` input was generated from running a small `jpeg` image through a `jpeg` reduction service<sup>4</sup>. The `libxml` input was selected from the `libxml`<sup>5</sup> test inputs smaller than 700 bytes with the greatest AFL branch coverage.

#### 5.2.1 Dataflow Accuracy

We evaluate the accuracy of PGA in comparison to DTA against an estimate of ground truth dataflows. This comparison setting favors DTA since it does not take the fine grained dataflow information from PGA into account (i.e., only considers binary 0/1 influence), but still illustrates the benefits of PGA’s increased precision. In addition to comparing against `dfsan`, we also compare against `libdft`, another widely used DTA framework that uses Intel PIN to instrument the binary directly, `Neutaint`, which uses the gradients of a neural network to model dataflows, and an ablation of PGA with binary gradients, `grsan` (`binary`). Notably, `libdft` tracks taint at byte level granularity and incorporates special case rules to handle operations that cancel out dataflows, such as  $y = x - x$ .

**Ground truth estimation.** To estimate ground truth dataflows, we measure if changes in taint sources cause changes in sink values during execution. When recording executions, we only consider executions that follow the same path to remove implicit flows, since neither DTA nor PGA can detect these. We mark each byte read from the input file as a source and each branch condition as a sink, because branches ultimately determine the behavior of a program, and because many security vulnerabilities can

<sup>3</sup><https://github.com/google/AFL>

<sup>4</sup><https://tinyjpg.com/>

<sup>5</sup><https://gitlab.gnome.org/GNOME/libxml2/>



	Neutaint			libdft			dfsan			grsan (binary)			grsan (floats)		
	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
minigzip	0.02	0.55	0.04	0.42	0.29	0.17	0.29	0.60	0.39	0.41	0.15	0.22	0.63	0.51	<b>0.57</b>
djpeg	0.02	0.33	0.04	-	-	-	0.22	1.00	0.37	0.62	0.63	0.62	0.60	0.83	<b>0.69</b>
mutool	0.002	0.19	0.004	0.70	0.32	0.22	0.63	0.61	0.62	0.87	0.50	0.63	0.86	0.51	<b>0.64</b>
xmllint	0.07	0.69	0.12	-	-	-	0.62	0.99	0.76	0.91	0.87	0.89	0.94	0.91	<b>0.92</b>
objdump	0.03	0.20	0.05	0.47	0.67	0.28	0.37	0.93	0.52	0.51	0.66	0.58	0.66	0.77	<b>0.71</b>
strip	0.02	0.39	0.03	0.26	0.59	0.18	0.20	0.96	0.33	0.42	0.72	0.53	0.50	0.86	<b>0.63</b>
size	0.06	0.39	0.11	0.20	0.59	0.30	0.37	0.95	0.53	0.54	0.76	0.63	0.62	0.91	<b>0.74</b>

Table 2: **Summary of accuracy comparison results for DTA and PGA systems. Neutaint, libdft, and dfsan are state-of-the-art DTA systems, while binary grsan is an ablation of PGA that only uses binary (1 or 0) gradients to test the impact of precise gradients on accuracy. Best F1 scores for each program are highlighted. Experiments with libdft on djpeg and xmllint timed out after 24hrs. PGA (with floating point gradients) outperforms DTA on all programs, and full precision (floats) grsan outperforms binary grsan on all programs.**

only be exploited when certain branches are taken. For each input byte, we set the byte to 0, 255, and toggling each bit for a total of 10 samples. We found that this sampling strategy usually triggered a change in the sink variable when there was a valid dataflow.

**Accuracy evaluation.** We perform the accuracy evaluation on the programs shown in Table 1 using a set of small seed files (<1Kb) to make sampling each byte feasible. Since valid dataflows often only involve a few input bytes, we use F1 accuracy, which is a standard metric for evaluating predictions on imbalanced classes in classification problems. F1 accuracy is computed as  $F1 = 2 * \frac{\text{precision} * \text{recall}}{\text{precision} + \text{recall}}$ . Precision indicates the proportion of bytes with predicted dataflows that are correct (i.e. not false positives), while recall indicates the proportion of valid dataflows that were correctly predicted (i.e. not false negatives). Results are shown in Table 2.

Generally, PGA achieves a significant improvement in precision, achieving up a 37% increase in precision and 33% increase in F1 accuracy (20% on average) compared to the best performing DTA system, `dfsan`. Overall PGA gets higher F1 scores for all programs. In spite of incorporating special case dataflow cancellation rules for its bitwise and numerical operations, `libdft` achieves lower accuracy than `dfsan` in the evaluation. We hypothesize this is due to the difficulty in writing handcrafted rules for all possible X86 instructions, which leads to errors in propagation rules as noted in [12]. The binary gradient PGA ablation, `grsan (binary)`, also has much lower accuracy than full precision PGA, indicating gradients are essential to computing accurate dataflows with PGA. We discuss the binary gradient ablation in more detail in Appendix A.

**Result 1:** PGA achieves the highest F1 accuracy on all 7 tested programs compared to 3 state-of-the-art DTA systems, and is up to 33% more accurate than the next most accurate DTA system, `dfsan`.

**Additional Accuracy Experiments.** In addition to the accuracy experiment in Table 2, we run experiments to address the following: (1) How do varying compiler optimization levels effect the accuracy of PGA vs. DTA? (2) How does PGA perform against Neutaint in Hot-byte prediction? (3) On which specific operations does PGA vary from DTA due to 0 gradients? (4) How does PGA compare with Quantitative Information Flow (QIF) techniques? We summarize the results here and describe these experiments in detail in Appendix A.

- 1. Compiler Optimization.** PGA’s accuracy improvement over DTA is robust to varying compiler optimization levels. On average, PGA is at least 18% more accurate than DTA on compiler optimization levels -O0 through -O2.
- 2. Hotbyte Prediction.** When we reproduce the Hot-byte experiment described in Neutaint [43], (i.e. identifying input bytes with the most dataflows to branches) PGA achieves 43.8% accuracy while Neutaint achieves 64.3% accuracy on average. Neutaint achieves higher average accuracy because it trains on a large corpus of recorded execution traces, while PGA and the DTA reason about a single input and execution trace at a time. We see Neutaint as a complementary method that performs well in identifying hotbytes, while PGA has better fine grained dataflow accuracy, and both methods could be used together in program analysis.
- 3. Zero Gradient Analysis.** PGA avoids overtainting errors when it computes zero gradients on operations where DTA would propagate taint. We find the zero gradients occur most frequently on And, Remainder, Sub, Mul, and Shift operations, and that zero gradients are most often caused by masking, shifting, or composition effects.
- 4. QIF Comparison.** We compare PGA with a QIF tool Flowcheck that quantifies information flow

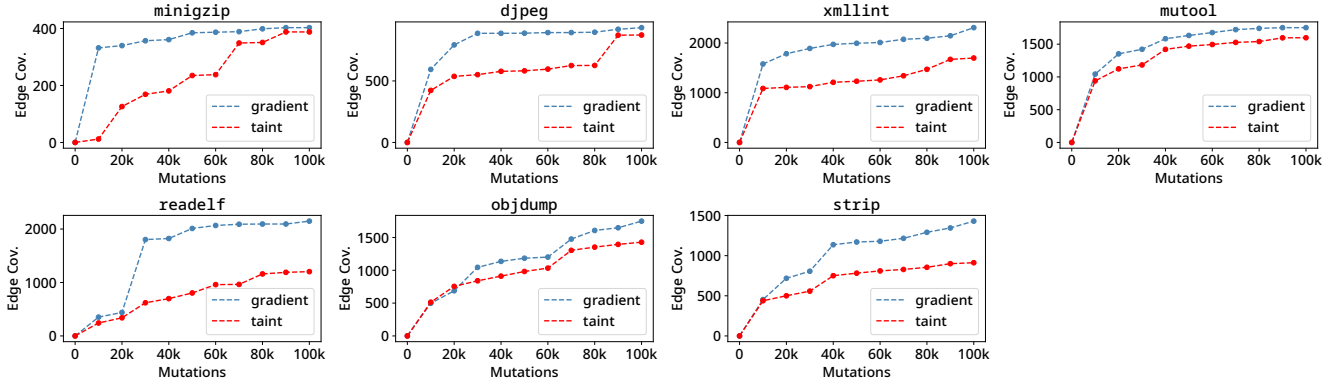


Figure 7: Comparison of guided fuzzer edge coverage achieved by PGA and DTA over 100k mutations from a single seed. Overall gradient-guided fuzzing achieves up to 56% higher coverage and improves the rate of new edge discovery by 10% on average.

in the form of bit leakage [28]. PGA outperforms Flowcheck by 22% on average in terms of F1 accuracy.

### 5.2.2 Overhead

We observe two conflicting phenomena when measuring overhead: PGA can either increase overhead due to the additional floating point storage and computation required by gradients, or decrease runtime and memory overhead when its increased precision reduces unnecessary dataflow tracking operations that use additional computation and shadow memory.

We evaluate the overhead introduced by our implementation of PGA in runtime and memory relative to `dfsan` on a single source dataflow. Note that if we consider overhead for multiple sources, the runtime will be lower and the memory overhead will be higher for a multi-source implementation. In the worst case PGA has 21.7% greater overhead in runtime and 21.5% in memory relative to DTA, but on average only adds 3.21% relative overhead in runtime and 1.48% in memory. Table 9 and Table 10 in Appendix B show the detailed results. We also provide overhead measurements for `libdft`, although it adds significantly more overhead due to the binary instrumentation.

**Result 2:** On average PGA increases runtime overhead by 3.21% runtime and memory overhead by 1.48% relative to DTA, and increases runtime by 21.7% and memory usage by 21.5% relative to DTA in the worst case.

### 5.2.3 Dataflow-Guided Fuzzing

Since dynamic dataflow analysis is often used as a tool to guide fuzzing, we evaluate PGA in comparison to DTA as a method for guiding fuzzer mutations. Unlike

our evaluation of dataflow accuracy, this experiment emphasizes the dataflow magnitude information provided by the program gradient, since bytes with the largest derivatives are selected for fuzzing.

We first compare PGA and DTA using a simple deterministic strategy for mutating input bytes based on dataflows to branches. This ensures there is no bias from randomized mutation strategies or other heuristics employed by state-of-the-art fuzzers in this evaluation. First, we execute the program with all inputs set as sources and all branches set as sinks. We then select 128 bytes from the input bytes based on the measured taint and gradient flows to branches. With PGA, the bytes with the greatest gradients are prioritized, this approach utilizes the additional information provided by PGA to improve the mutation strategy. The fuzzer performs a deterministic set of mutations on the selected 128 bytes, in which each byte in turn is set to all 256 possible values.

**Edge coverage comparison.** We execute the fuzzer with both PGA and DTA for 100,000 mutations, and record coverage every 10,000 mutations. Figure 7 shows the relative edge coverage achieved by each method. On average the gradient guided fuzzing outperforms taint in increasing edge coverage by 10% per 10,000 mutations. The gradient guided fuzzer achieves higher coverage on all programs, with the greatest improvement in overall edge coverage of 56% on `strip`. We note that for some programs such as `xmllint`, there is a significant disparity between the results of the guided fuzzing and precision evaluations. We believe this difference is caused by two factors: the magnitude of the gradient was more important than its accuracy in guiding the fuzzer on these programs, and that even small differences in accuracy can be significant if they allow the fuzzer to precisely target key branches in the program.

**Enhancing state-of-the-art fuzzers.** We also evalu-

ate if the gradient information from PGA can improve the performance of NEUZZ, a state-of-the-art fuzzer. We evaluate a version of NEUZZ modified to use PGA against unmodified NEUZZ and VUZZer, another dataflow guided fuzzer. On average, PGA+NEUZZ improves new edge coverage by 12.9% over baseline NEUZZ. We hypothesize this improvement is because the gradients produced with PGA are more precise than the neural-network based gradients used by NEUZZ. We discuss this experiment and provide more detailed results in Appendix C.

**Result 3:** In guided fuzzing PGA increases the rate of edge coverage growth by 10% on average compared to DTA, and improves the edge coverage of NEUZZ, a state-of-the-art fuzzer, by 12.9% on average.

### 5.3 Bug Finding

Next, we show the additional information provided by PGA make it a useful tool for discovering and analyzing different types of bugs in real world programs. We test PGA against DTA in three applications: detecting and analyzing known vulnerabilities, guiding discovery of new bugs, and discovering information leaks.

#### 5.3.1 Analysis of known CVEs

We first evaluate PGA as a tool for detecting dangerous dataflows offline in known CVEs. We instrument the programs to mark user-controlled input as dataflow sources and the instructions involved in the attacks as dataflow sinks. We select 21 CVEs that cover a range of vulnerability types, including stack and heap overflows, integer overflows, memory allocation errors, and null pointer dereferences. We include CVEs from our evaluation programs as well as `openssl` to demonstrate PGA based analysis on a variety of program types.

Table 3 shows a comparison of PGA and DTA in detecting the relevant dataflows in these CVEs. PGA correctly identifies dataflows for 19 out of the 21 evaluated CVEs, including 2 CVEs that cannot be identified with DTA. For these CVEs, DTA overtaints on the malicious inputs and crashes due to label exhaustion, while PGA can precisely identify the dataflows without over-tainting. For the 2 CVEs which both PGA and DTA fail to detect, the dataflow source indirectly propagates to the sink through implicit dataflows (i.e. control flow).

We also note the utility of the additional information provided by gradients and how it can help distinguish vulnerabilities in an online manner. In the case of CVE-2017-15996, an out of memory allocation error triggered by the dataflow from an input byte, PGA directly measures the effect of input changes on the size of the allocation, and can early terminate when it finds input byte values that will trigger the out-of-memory error.

CVE ID	Vulnerability - Program	PGA DTA	
CVE-2007-1657	stack overflow - minigzip	✓	✓
CVE-2017-7210	off-by-one read - objdump	✓	✓
CVE-2017-8396	heap overflow - libbfd	✓	✓
CVE-2017-15996	out-of-memory - readelf	✓	✓
CVE-2018-6543	integer overflow - objdump	✓	✓
CVE-2018-6759	null ptr dereference - nm	✓	✓
CVE-2018-7643	integer overflow - objdump	✓	✓
CVE-2018-10372	heap overflow - readelf	✓	✓
CVE-2018-11813	infinite loop - cjpeg	✓	✓
CVE-2018-12698	out-of-memory - libiberty	✓	✓
CVE-2018-12699	heap overflow - libiberty	✓	✓
CVE-2020-14152	out-of-memory - djpeg	✓	✓
CVE-2018-19932	integer overflow - strip	✓	✓
CVE-2018-19777	infinite loop - mutool	✓	✓
CVE-2018-20671	infinite loop - objdump	✓	✓
CVE-2019-14444	integer overflow - readelf	✓	✓
CVE-2020-1967	null ptr dereference - openssl	✓	✓
CVE-2018-11212	divide-by-zero - cjpeg	✓	×
CVE-2018-11214	heap overflow - cjpeg	✓	×
CVE-2020-7041	invalid certificate - openssl	×	×
CVE-2018-12697	null ptr dereference - libiberty	×	×

Table 3: List of 21 CVEs for which the exploitable dataflows were analyzed by PGA and DTA (`dfsan` and `libdft`)

**Result 4:** PGA identifies relevant dataflows in 19 out of 21 evaluated CVEs, including 2 DTA cannot detect due to label exhaustion. PGA and DTA both cannot identify control-flow-based dataflows for 2 CVEs.

#### 5.3.2 Bug Discovery

We compare PGA and DTA as bug discovery tools by adding additional instrumentation to record dataflows for instruction and function arguments that can potentially trigger program errors, such as memory allocations, copy instructions, indexing operations, and shift operators. We then execute the programs on a corpus of files generated by running AFL on each program for 24 hours. Next, we generate new inputs by changing input bytes involved in the recorded dataflows similar to Section 5.2.3. For PGA, we select 128 input bytes prioritized based on the function gradient, while for DTA, we randomly select them. We modify the values of the selected bytes based on the gradient for PGA or by setting them to 0 or 255 for DTA.

Table 4 summarizes our results. Overall, PGA finds 22 bugs in our evaluated programs through gradient guided modification of the inputs, including arithmetic errors, out-of-memory allocations, and integer overflows. The DTA guided bug search finds 15 of these 22 bugs. Of the 22 bugs, 20 have been confirmed by the developers, 3 of them resulted in new patches, and the remaining 17 were already patched in the latest sources of the programs.

```

1  GRSAN_MARK_BYTE(c, 1.0); //grad = 1.0
2
3  cinfo->A1 = (c) & 15; //grad = 1.0
4  ...
5  (*block)[natural_order[k]] =
6      (JCOEF) (v << cinfo->A1);
7  /* block[0] gradient = 8.0 */
8
9  void jpeg_idct_islow(int * block) {
10     ...
11     int * inptr = block; // grad = 8.0
12
13     z2 = (int) inptr[0] * quantptr[0]
14     /*z2 gradient = 2040.0 can overflow*/
15
16     z2 = z2 << 13;
17     /* negative z2 triggers error */
18 }

```

Figure 8: Arithmetic Error in `djpeg`.

Library	Test Program	Integer Overflow	Memory Corruption
libjpeg-9c	<code>djpeg</code>	2	3
mupdf-1.14.0	<code>mutool show</code>	1	0
binutils-2.30	<code>size</code>	0	1
	<code>objdump -xD</code>	0	9
	<code>strip</code>	0	6

Table 4: Summary of new bugs found by PGA. In total there are 22 bugs found over 5 programs.

For the 7 bugs that were found by PGA and not DTA, gradient magnitude and direction allowed the search to prioritize input bytes that could trigger errors that could not be identified with DTA. We give a case study in Figure 8, which illustrates how large gradients are used to find an arithmetic error in `djpeg`. By altering an input byte with a large gradient to a shift operand, an overflow is triggered that results in an invalid operation. Similarly, identifying inputs with large gradients to memory operations was key to finding memory errors.

**Result 5:** A simple PGA guided search finds 22 bugs in the tested programs. A DTA guided search using the same strategy and inputs finds 15 of these 22 bugs.

### 5.3.3 Information Leak Discovery

We provide two case studies using PGA to detect side channel leaks: one example of a memory usage based side channel in `objdump` and an execution time based side channel in `cjpeg`. To identify each information leak, we marked the input file headers as sources and relevant program values as sinks, either memory allocation operands or comparison operands in loops.

In `objdump`, we identified a memory based side channel based on a gradient of 1 million to a `malloc` instruction from the ELF section header for program size. Figure 9a shows the effect of incrementing the value from 46 to

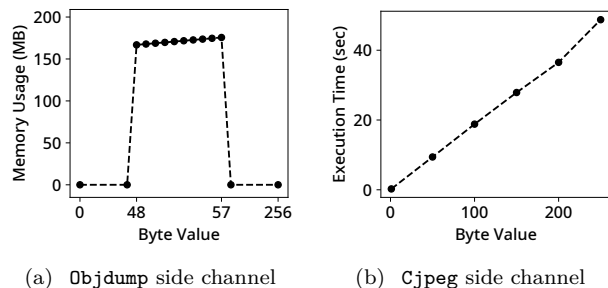


Figure 9: Memory and timing side channel leaks.

59 on the program’s total memory usage. The memory consumption is linear in the byte value if the byte is in range from 48 to 57, which can be converted to a valid number ‘0’ to ‘9’ in ASCII. Similarly, we identified the timing based side-channel in `cjpeg` by a gradient from the height field in the jpeg header to the operand of a while loop condition. Figure 9b shows the height information leak in program execution time.

Prior side channel attacks have demonstrated that these types of leaks can be exploited to learn sensitive information about a user [11, 24]. For example, one can imagine a malicious Android app that uses JPEG dimensions leaked from a browser to determine which websites the device user is visiting.

**Result 6:** PGA successfully detects two information leaks from file headers in `objdump` and `cjpeg`.

## 6 Discussion

In this section we review the implications of our results and discuss of the relative advantages and limitations of PGA as an approach to dataflow analysis.

**Advantages of PGA.** The additional information encoded in gradients can greatly improve precision when predicting dataflows between sources and sinks (i.e. reducing the number of false positives), while the magnitude and direction information can be used to prioritize dataflows based on their significance and predicted effect. We see the benefits of the additional information from gradients in the improved performance of PGA relative to DTA in our dataflow accuracy, guided fuzzing, and vulnerability detection and analysis evaluations (Sections 5.2.1, 5.2.3, 5.3).

**Limitations of PGA.** While our implementation of PGA demonstrably works based on our evaluation, How to best sample non-smooth operations when evaluating proximal gradients is an open question. Our prototype uses a simple fixed sampling strategy, and does not fully implement proximal gradients on some operations, such as loads on pointers with derivatives (Section 3.3), which sometimes causes errors in the gradient evaluation. The



effect of these errors can be seen in our dataflow accuracy evaluation (Section 5.2.1), where `grsan` has slightly lower recall than `dfsan`, indicating some gradients erroneously evaluate to 0. We believe incorporating more information about specific operations in sampling strategies, as well as tracking valid domains for some operations, will reduce these errors.

A second limitation of our implementation of PGA is that, like most DTA frameworks, it does not model implicit dataflows, such as control flow dependencies. This can be seen in our CVE evaluation (Section 5.3.1), where two of our tested CVEs cannot be detected by either PGA or DTA. We intend to explore both more accurate methods for evaluating proximal gradients and modeling implicit dataflows in future work.

## 7 Related Work

**Dynamic Taint Analysis.** Dynamic Taint Analysis (DTA) tracks data flow from taint sources to taint sinks at runtime. Common applications of DTA include software vulnerability analysis and information leak detection [15, 18, 34, 58, 59]. DTA typically overestimates the tainted bytes which contributes to a large performance overhead. Therefore, much of the recent work in DTA has focused on developing more efficient systems [7, 25, 29]. Like DTA, PGA dynamically propagates dataflow information, but it provides more fine-grained information in the form of gradients. Moreover, PGA is more precise than DTA, which reduces overtainting in large programs.

Some DTA systems use bit level taint tracking to improve precision at the cost of higher overheads [54, 55]. Although we have not implemented it in our current prototype, gradients can also be propagated over individual bits based on functional Boolean analysis, and we expect it to offer similar tradeoffs in improved accuracy for higher overheads [35].

Recently, automatically learning taint rules has been used to reduce the approximation errors in DTA [12]. This approach is orthogonal to ours and could also potentially be applied to learn gradient propagation rules.

**Quantitative Information Flow.** Quantitative Information Flow (QIF) measures the potential transmission of information through a program using entropy based measures such as channel capacity and min-entropy [19, 28, 50]. QIF has primarily been used for detecting information leaks and ensuring the integrity of program secrets [3, 17, 23], but has also been proposed as a way of enhancing taint tracking [33]. PGA adds a different type of information as discussed in Appendix A, and does not have the high computational complexity involved in estimating information flows accurately.

**Gradient-guided fuzzing.** Recent fuzzers have used gradient approximations to guide their mutation process.

Angora estimates finite differences, an approximation of gradients with many known limitations especially for high-dimensional problems, by executing the program on modified inputs and recording the changes in the outputs [10, 39]. `NEUZZ`, `MTFuzz` and `Neutaint` train neural networks to predict program branch behavior and use the network’s gradients to guide the mutation algorithm [43–45]. This incurs less overhead than instrumentation based methods but is also less exact since it operates on an approximate model of the program. By contrast, PGA computes gradients directly over the program’s individual instructions and therefore produces precise gradients.

**Program Smoothing.** Prior work has explored computing gradients with smooth interpretation of a program via a Gaussian kernel [8, 9] or parametric relaxation of SMT [41, 57]. These methods use symbolic reasoning and have not been applied to analysis of real world programs. PGA’s approximation methods are more efficient and have been successfully demonstrated on real world programs.

## 8 Conclusion

In this paper we introduce proximal gradient analysis (PGA), a novel theoretically-grounded approach to dataflow analysis that uses non-smooth calculus techniques to compute gradients over programs. PGA is more precise than dynamic taint tracking and provides more fine grained information about program behavior. We provide a prototype implementation of PGA based on the LLVM framework and show that it outperforms three state-of-the-art DTA systems in accuracy while adding less than 5% overhead on average. Finally, we show PGA is an effective tool for security analysis, identifying relevant dataflows for 19 different CVEs, discovering 22 bugs, and detecting 2 side-channel leaks in 7 real world programs. We hope that our approach to program analysis will motivate other researchers to explore new techniques exploiting the rich non-smooth analysis literature.

## Acknowledgements

We thank our shepherd Lujo Bauer and the anonymous reviewers for their constructive and valuable feedback. The first author is supported by an NDSEG Fellowship, and the second author is supported by an NSF Graduate Fellowship. This work is sponsored in part by NSF grants CNS-18-42456, CNS-18-01426, CNS-16-17670; ONR grant N00014-17-1-2010; an ARL Young Investigator (YIP) award; a NSF CAREER award; a Google Faculty Fellowship; and a Capital One Research Grant, as well as European Union Marie Skłodowska-

Curie grant agreement 690972 (PROTASIS) and innovation programme under grant agreement No. 786669 (ReAct). Any opinions, findings, conclusions, or recommendations expressed herein are those of the authors, and do not necessarily reflect those of the US Government, European Union, ONR, ARL, NSF, Google, or Capital One.

## References

- [1] Pin - A Dynamic Binary Instrumentation Tool. <https://software.intel.com/en-us/articles/pin-a-dynamic-binary-instrumentation-tool>, 2019.
- [2] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, and Zhifeng Chen. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [3] Mário S Alvim, Konstantinos Chatzikokolakis, Annabelle McIver, Carroll Morgan, Catuscia Palamidessi, and Geoffrey Smith. Additive and multiplicative notions of leakage, and their capacities. In *2014 IEEE 27th Computer Security Foundations Symposium*, pages 308–322. IEEE, 2014.
- [4] Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel, Jacques Klein, Yves Le Traon, Damien Octeau, and Patrick McDaniel. Flowdroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for android apps. *Acm Sigplan Notices*, 49(6):259–269, 2014.
- [5] David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Müller. How to explain individual classification decisions. *Journal of Machine Learning Research*, 11:1803–1831, 2010.
- [6] Davide Balzarotti, Marco Cova, Vika Felmetzger, Nenad Jovanovic, Engin Kirda, Christopher Kruegel, and Giovanni Vigna. Saner: Composing static and dynamic analysis to validate sanitization in web applications. In *IEEE Symposium on Security and Privacy (S&P)*, pages 387–401. IEEE, 2008.
- [7] Erik Bosman, Asia Slowinska, and Herbert Bos. Minemu: The world’s fastest taint tracker. In *RAID*, 2011.
- [8] Swarat Chaudhuri and Armando Solar-Lezama. Smooth interpretation. *ACM Sigplan Notices*, 45(6):279–291, 2010.
- [9] Swarat Chaudhuri and Armando Solar-Lezama. Smoothing a program soundly and robustly. In *Computer Aided Verification*, pages 277–292, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [10] Peng Chen and Hao Chen. Angora: Efficient fuzzing by principled search. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 711–725. IEEE, 2018.
- [11] Shuo Chen, Rui Wang, XiaoFeng Wang, and Kehuan Zhang. Side-channel leaks in web applications: A reality today, a challenge tomorrow. In *IEEE Symposium on Security and Privacy*, pages 191–206. IEEE, 2010.
- [12] Zheng Leong Chua, Yanhao Wang, Teodora Baluta, Prateek Saxena, Zhenkai Liang, and Purui Su. One engine to serve’em all: Inferring taint rules without architectural semantics. 2019.
- [13] Frank H Clarke. *Optimization and nonsmooth analysis*, volume 5. Siam, 1990.
- [14] James Clause, Wanchun Li, and Alessandro Orso. Dytan: a generic dynamic taint analysis framework. In *International Symposium on Software Testing and Analysis*, pages 196–206. ACM, 2007.
- [15] James Clause, Wanchun Li, and Alessandro Orso. Dytan: A generic dynamic taint analysis framework. 2007.
- [16] R Dennis Cook. Assessment of local influence. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 133–169, 1986.
- [17] Goran Doychev, Boris Köpf, Laurent Mauborgne, and Jan Reineke. Cacheaudit: A tool for the static analysis of cache side channels. *ACM Transactions on Information and System Security (TISSEC)*, 18(1):4, 2015.
- [18] William Enck, Peter Gilbert, Byung-Gon Chun, Landon P. Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N. Sheth. Taintdroid: An information-flow tracking system for realtime privacy monitoring on smartphones. In *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation, OSDI’10*, 2010.
- [19] Barbara Espinoza and Geoffrey Smith. Min-entropy as a resource. *Information and Computation*, 226:57–75, 2013.
- [20] Vijay Ganesh, Tim Leek, and Martin Rinard. Taint-based directed whitebox fuzzing. In *Proceedings of*

- the 31st International Conference on Software Engineering*, pages 474–484. IEEE Computer Society, 2009.
- [21] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [22] Andreas Griewank. Automatic directional differentiation of nonsmooth composite functions. In *Recent Developments in Optimization*, pages 155–169. Springer, 1995.
- [23] Jonathan Heusser and Pasquale Malacaria. Quantifying information leaks in software. In *Proceedings of the 26th Annual Computer Security Applications Conference*, pages 261–269. ACM, 2010.
- [24] Suman Jana and Vitaly Shmatikov. Memento: Learning secrets from process footprints. In *IEEE Symposium on Security and Privacy*, pages 143–157. IEEE, 2012.
- [25] Vasileios P. Kemerlis, Georgios Portokalidis, Kangkook Jee, and Angelos D. Keromytis. Libdft: Practical dynamic data flow tracking for commodity systems. In *Proceedings of the 8th ACM SIGPLAN/SIGOPS Conference on Virtual Execution Environments*, VEE '12, 2012.
- [26] Boris Köpf and Andrey Rybalchenko. Approximation and randomization for quantitative information-flow analysis. In *2010 23rd IEEE Computer Security Foundations Symposium*, pages 3–14. IEEE, 2010.
- [27] Chris Lattner and Vikram Adve. Llvm: A compilation framework for lifelong program analysis & transformation. In *Proceedings of the International Symposium on Code Generation and Optimization: Feedback-directed and Runtime Optimization*, CGO '04, pages 75–, Washington, DC, USA, 2004. IEEE Computer Society.
- [28] Stephen McCamant and Michael D Ernst. Quantitative information flow as network flow capacity. In *ACM SIGPLAN Notices*, volume 43, pages 193–205. ACM, 2008.
- [29] Jiang Ming, Dinghao Wu, Jun Wang, Gaoyao Xiao, and Peng Liu. Straighttaint: Decoupled offline symbolic taint analysis. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering*, ASE, New York, NY, USA, 2016. ACM.
- [30] Andrew C. Myers and Andrew C. Myers. Jflow: Practical mostly-static information flow control. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '99, 1999.
- [31] Yu Nesterov. Lexicographic differentiation of nonsmooth functions. *Mathematical programming*, 104(2-3):669–700, 2005.
- [32] Nicholas Nethercote and Julian Seward. Valgrind: a framework for heavyweight dynamic binary instrumentation. *ACM Sigplan notices*, 42(6):89–100, 2007.
- [33] James Newsome, Stephen McCamant, and Dawn Song. Measuring channel capacity to distinguish undue influence. In *Proceedings of the ACM SIGPLAN Fourth Workshop on Programming Languages and Analysis for Security*, pages 73–85. ACM, 2009.
- [34] James Newsome and Dawn Xiaodong Song. Dynamic taint analysis for automatic detection, analysis, and signature generation of exploits on commodity software. In *NDSS*, volume 5, pages 3–4. Citeseer, 2005.
- [35] Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [36] Neal Parikh, Stephen Boyd, et al. Proximal algorithms. *Foundations and Trends in Optimization*, 1(3):127–239, 2014.
- [37] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. Deepxplore: Automated whitebox testing of deep learning systems. In *proceedings of the 26th Symposium on Operating Systems Principles*, pages 1–18. ACM, 2017.
- [38] Sanjay Rawat, Vivek Jain, Ashish Kumar, Lucian Cojocar, Cristiano Giuffrida, and Herbert Bos. VUzzer: Application-Aware Evolutionary Fuzzing. In *Proceedings of the 2008 Network and Distributed Systems Security Conference*, 2017.
- [39] D. Richtmeyer and K.W. Morton. *Difference Methods for Initial Value Problems*. Wiley, 2nd edition, 1967.
- [40] R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*, volume 317. Springer Science & Business Media, 2009.
- [41] Gabriel Ryan, Justin Wong, Jianan Yao, Ronghui Gu, and Suman Jana. Cln2inv: Learning loop invariants with continuous logic networks. In *International Conference on Learning Representations*, 2019.

- [42] Edward J Schwartz, Thanassis Avgerinos, and David Brumley. All you ever wanted to know about dynamic taint analysis and forward symbolic execution (but might have been afraid to ask). In *IEEE Symposium on Security and Privacy (S&P)*, pages 317–331. IEEE, 2010.
- [43] Dongdong She, Yizheng Chen, Baishakhi Ray, and Suman Jana. Neutaint: Efficient dynamic taint analysis with neural networks. *arXiv preprint arXiv:1907.03756*, 2019.
- [44] Dongdong She, Rahul Krishna, Lu Yan, Suman Jana, and Baishakhi Ray. Mtfuzz: Fuzzing with a multi-task neural network. *Foundations of Software Engineering*, 2020.
- [45] Dongdong She, Kexin Pei, Dave Epstein, Junfeng Yang, Baishakhi Ray, and Suman Jana. NEUZZ: Efficient Fuzzing with Neural Program Smoothing. In *IEEE Symposium on Security and Privacy (S&P)*, 2019.
- [46] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. *arXiv preprint arXiv:1704.02685*, 2017.
- [47] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [48] Asia Slowinska and Herbert Bos. Pointless tainting? evaluating the practicality of pointer tainting. In *Proceedings of the 4th ACM European conference on Computer systems*, pages 61–74. ACM, 2009.
- [49] Asia Slowinska and Herbert Bos. Pointer tainting still pointless:(but we all see the point of tainting). *ACM SIGOPS Operating Systems Review*, 44(3):88–92, 2010.
- [50] Geoffrey Smith. Quantifying information flow using min-entropy. In *2011 Eighth International Conference on Quantitative Evaluation of SysTems*, pages 159–167. IEEE, 2011.
- [51] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. Deeptest: Automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th international conference on software engineering*, pages 303–314, 2018.
- [52] Doug Ward. Chain rules for nonsmooth functions. *Journal of Mathematical Analysis and Applications*, 158(2):519–538, 1991.
- [53] R. E. Wengert. A simple automatic derivative evaluation program. *Commun. ACM*, 7(8):463–464, August 1964.
- [54] Babak Yadegari and Saumya Debray. Bit-level taint analysis. In *2014 IEEE 14th International Working Conference on Source Code Analysis and Manipulation*, pages 255–264. IEEE, 2014.
- [55] LK Yan, A Henderson, X Hu, H Yin, and S McCamant. On soundness and precision of dynamic taint analysis. *Dep. Elect. Eng. Comput. Sci., Syracuse Univ., Tech. Rep. SYR-EECS-2014-04*, 2014.
- [56] Lok Kwong Yan and Heng Yin. Droidscape: Seamlessly reconstructing the os and dalvik semantic views for dynamic android malware analysis. In *USENIX Security*, pages 569–584, 2012.
- [57] Jianan Yao, Gabriel Ryan, Justin Wong, Suman Jana, and Ronghui Gu. Learning nonlinear loop invariants with gated continuous logic networks. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 106–120, 2020.
- [58] Heng Yin, Dawn Song, Manuel Egele, Christopher Kruegel, and Engin Kirda. Panorama: Capturing system-wide information flow for malware detection and analysis. In *Proceedings of the 14th ACM Conferences on Computer and Communication Security (CCS)*, 2007.
- [59] Insu Yun, Sangho Lee, Meng Xu, Yeongjin Jang, and Taesoo Kim. QSYM: A Practical Concolic Execution Engine Tailored for Hybrid Fuzzing. In *Proceedings of the 27th USENIX Security Symposium (Security)*, Baltimore, MD, August 2018.

	Opt 0		Opt 1		Opt 2	
	dfsan	grsan	dfsan	grsan	dfsan	grsan
minigzip	0.39	0.57	0.42	0.51	0.39	0.45
djpeg	0.36	0.69	0.29	0.65	0.31	0.63
mutool	0.62	0.64	0.56	0.66	0.52	0.62
xmllint	0.76	0.92	0.73	0.88	0.74	0.81
objdump	0.52	0.71	0.48	0.67	0.47	0.68
strip	0.33	0.63	0.31	0.60	0.31	0.61
size	0.53	0.74	0.52	0.68	0.51	0.69

Table 5: **Effect compiler optimization levels on dataflow F1 accuracy. The table shows grsan has significantly higher F1 accuracy than dfsan for all three measured optimization levels (> 18% average)**



## A Additional Accuracy Experiments

We describe the additional accuracy evaluations summarized in Section 5.2.1 here. Specifically, the ablation of gradient information, the effects of compiler optimization, PGA vs. `Neutaint` in coarse grained dataflow prediction, analysis of 0 gradients, and a comparison with QIF.

**Gradient Ablation.** We measure the effect of the gradient information on determining accurate dataflows by performing an ablation with binary valued gradients. The ablation uses the same proximal gradient propagation rules, but rounds all gradients to 0 or 1. In effect, this converts PGA into DTA with PGA propagation.

Results of the comparison are shown in Table 2. PGA with floating point gradient information performs significantly better than PGA with binary gradients for every program. These results indicate that precise gradients are key to the performance gains achieved by PGA because they compose accurately over multiple operations.

**Compiler Optimization.** We evaluate the impact of compiler optimization levels on dataflow accuracy at 3 optimization levels: `-O0`, `-O1`, and `-O2`. Table 5 summarizes the effects of 3 compiler optimization levels on dataflow F1 accuracy. Increasing the compiler optimization levels reduces the accuracy of both PGA and DTA by a small amount (<3.6%) for both `-O1` and `-O2`. On average, PGA is at least 18% more accurate than DTA for all three tested optimization levels.

**Neutaint hotbyte evaluation.** `Neutaint`'s neural network based approach does not perform well in fine grained dataflow prediction, but is better suited to identifying hot bytes (input bytes that are most influential to program behavior). We therefore perform the hot byte evaluation described in [43] on PGA. Our results are summarized in Table 6. On average, PGA predicts hotbytes with 43.75% accuracy, while `Neutaint` predicts hotbytes with 64.25% accuracy. We see `Neutaint` as a complementary method to PGA, where PGA is better suited to fine grained dataflow prediction and both methods could be used together in program analysis.

**Zero gradient analysis.** PGA is able to avoid over-tainting when it computes a zero gradient on an instruction DTA would mark as tainted. Therefore we investigate the distribution of zero gradients across programs and instruction types to determine where and how PGA is more precise than DTA. For each program and each type of instruction, we count how many times the instruction had zero gradient in the execution traces from the accuracy evaluation. Table 8 shows the results of this analysis for each instruction and program.

**QIF Comparison.** We compare PGA with the latest version of a publicly available QIF tool `Flowcheck` [28]. We perform a similar experiment to Section 5.2.1, but since `Flowcheck` does not byte-level granularity, we com-

Program	Neutaint	PGA
<code>mutool</code>	73%	99%
<code>xmllint</code>	76%	1%
<code>djpeg</code>	37%	33%
<code>miniunz</code>	71%	42%

Table 6: **Neutaint Hotbyte Evaluation results. On average, Neutaint predicts hot bytes with 64.25% accuracy and PGA with 43.75% accuracy. We believe Neutaint outperforms PGA because it makes a prediction based on many program inputs, whereas PGA makes a prediction based on a single input. Note our results different from the original Neutaint paper [43] due to different initializations and environments for training the neural network.**

	Flowcheck			PGA		
	Prec.	Rec.	F1	Prec.	Rec.	F1
<code>minigzip</code>	0.44	0.62	0.52	0.7	0.62	0.66
<code>djpeg</code>	0.44	0.95	0.6	0.62	0.87	0.73
<code>mutool</code>	0.69	0.77	0.73	0.89	0.64	0.74
<code>xmllint</code>	0.55	0.08	0.14	0.99	0.95	0.97
<code>objdump</code>	0.71	0.62	0.66	0.75	0.79	0.77
<code>strip</code>	0.49	0.65	0.56	0.66	0.88	0.75
<code>size</code>	0.74	0.64	0.69	0.74	0.93	0.82

Table 7: **QIF accuracy comparison results for PGA and Flowcheck. PGA outperforms Flowcheck by 22% on average in terms of F1 accuracy.**

pute accuracy by aggregating flows over all bytes so that PGA is not unfairly advantaged. We outperform `Flowcheck` in terms of F1 accuracy by 22% on average on all of the evaluated programs as summarized in Table 7.

## B Runtime and Memory Overhead Evaluation

**Program Overhead.** We evaluate the overhead introduced by our implementation of PGA in runtime and memory and compare it to `dfsan` for a single taint/gradient source. To measure overhead, we execute each program while recording runtime and memory usage. For runtime we perform 5,000 executions for each measurement. We perform each measurement 5 times and average the measured runtime and memory usage.

Tables 9 and 10 detail the runtime and memory overhead per program in our evaluation. In the worst case PGA has 21.7% greater overhead in runtime and 21.5% in memory relative to DTA, but on average only adds 3.21 % relative overhead in runtime and 1.48% in memory. We also provide overhead measurements for `libdft`, although it adds significantly more overhead due to the binary instrumentation.

Program Summary Over all Instructions			Instruction Summary Across all Programs		
Program	Instrs	%Zeros	Instr.	Total	%Zeros
minigzip	3012	28.2	And	6756	30.2
djpeg	703	38.7	URem	214	29.0
mutool	401	40.4	Sub	1214	21.0
xmllint	430	39.5	Mul	875	15.9
objdump	1070	39.0	LShr	2377	14.4
strip	3089	41.0	AShr	149	6.0
size	659	19.3	Add	895	5.7

Table 8: Analysis of operations from execution traces where gradient drops to 0, aggregated for each program and for each type of instruction across all programs. Outputs of these operations will have 0 gradient but still be marked as tainted by DTA.

Program	libdft Overhead	dfsan Overhead	grsan Overhead	grsan rel. to dfsan
minigzip	2,379.5%	54.7%	61.5%	4.4%
djpeg	-	70.5%	73.7%	1.9%
mupdf	853.5%	198.4%	262.1%	21.5%
xmllint	231.4%	5.5%	0.0%	-5.2%
size	152.5%	101.1%	107.1%	3.0%
objdump	180.0%	133.2%	131.2%	-0.9%
strip	142.5%	12.0%	11.4%	-2.2%

Table 9: Program runtime overhead measurements averaged over five runs for a single taint/gradient source. Libdft overhead is measured relative to running a program only with PIN. Dfsan and grsan are measured relative to uninstrumented programs. After 6 hours, libdft execution timed out on djpeg.

Program	dfsan Overhead	grsan Overhead	grsan rel. to dfsan
minigzip	183.7%	245.3%	21.7%
djpeg	276.4%	291.9%	4.1%
mupdf	112.4%	124.7%	5.8%
xmllint	346.6%	258.5%	-19.7%
size	373.3%	392.4%	4.0%
objdump	345.6%	323.5%	-5.0%
strip	344.5%	342.1%	-0.5%

Table 10: Memory overhead for each program averaged over five runs relative to uninstrumented programs for a single taint/gradient source. Grsan may increase or decrease overhead because gradients require more memory to store, but may use less overall memory due to increased precision. On average, grsan adds 1.48% additional overhead relative to dfsan.

Program	Edge Coverage after 24hrs			
	VUzzer	NEUZZ	PGA + NEUZZ	PGA+NEUZZ rel. to NEUZZ
minigzip	-	87	94	8.1%
djpeg	7	645	686	6.4%
mupdf	156	376	430	14.4%
xmllint	282	957	1079	12.8%
size	474	1580	2064	30.6%
objdump	247	1813	2014	11.1%
strip	1337	3394	3637	7.2%

Table 11: New edge coverage for each program over 24 hours by three different fuzzers. VUzzer encounters an error in its taint tracking on minigzip and crashes. Overall, PGA+NEUZZ improves NEUZZ edge coverage on average by 12.9%. Note that our results are slightly different from the original NEUZZ and VUzzer results due to differences in test environments, input corpuses, and program versions.

## C Evaluation on Current Fuzzers

We also evaluate if the gradient information from PGA can improve the performance of state-of-the-art fuzzers such as NEUZZ and VUzzer. We use NEUZZ as a basis because it has higher edge coverage as seen in Table 11 and already incorporates gradients from a neural network in its mutation strategy. We modify NEUZZ so that it uses the PGA gradients to guide its mutation strategy. We run grsan on its inputs and send the resulting gradients to the NEUZZ backend. Note that NEUZZ is designed to operate on gradients, so we did not modify it to also use DTA. We provide a controlled comparison of PGA vs. DTA for guided fuzzing in Evaluation 5.2.3.

We compare the additional edge coverage achieved by the fuzzers over a 24hr run. Since we use some programs with different file formats from the original NEUZZ benchmark, we use a new seed corpus generated by running AFL on each program for 1 hour. We perform this experiment using cloud hosted virtual machines. Table 11 summarizes the modified PGA+NEUZZ against baseline NEUZZ and VUzzer. On average, PGA+NEUZZ improves new edge coverage by 12.9% over baseline NEUZZ. We hypothesize this improvement is because the gradients produced with PGA are more precise than the neural-network based gradients used by NEUZZ. The very similar results in edge coverage on minigzip are caused by the CRC check in minigzip, which causes the program to exit early on most new inputs. VUzzer crashes on minigzip due to an error in its taint tracking and achieves a low edge coverage for djpeg because of the high overhead of PIN’s dynamic binary instrumentation for taint tracking. Note that our results are slightly different from the original NEUZZ and VUzzer results due to different initial seed corpuses, program versions, and test environments.