Automatic Input Rectification

Fan Long, Vijay Ganesh, Michael Carbin, Stelios Sidiroglou, and Martin Rinard
MIT CSAIL
{*fanl, vganesh, mcarbin, stelios, rinard*}*@csail.mit.edu*

***Abstract*—We present a novel technique, *automatic input rectification*, and a prototype implementation, SOAP. SOAP learns a set of constraints characterizing *typical inputs* that an application is highly likely to process correctly. When given an *atypical input* that does not satisfy these constraints, SOAP automatically *rectifies* the input (i.e., changes the input so that it satisfies the learned constraints). The goal is to automatically convert potentially dangerous inputs into typical inputs that the program is highly likely to process correctly.**

Our experimental results show that, for a set of benchmark applications (Google Picasa, ImageMagick, VLC, Swfdec, and Dillo), this approach effectively converts malicious inputs (which successfully exploit vulnerabilities in the application) into benign inputs that the application processes correctly. Moreover, a manual code analysis shows that, if an input does satisfy the learned constraints, it is incapable of exploiting these vulnerabilities.

We also present the results of a user study designed to evaluate the subjective perceptual quality of outputs from benign but atypical inputs that have been automatically rectified by SOAP to conform to the learned constraints. Specifically, we obtained benign inputs that violate learned constraints, used our input rectifier to obtain rectified inputs, then paid Amazon Mechanical Turk users to provide their subjective qualitative perception of the difference between the outputs from the original and rectified inputs. The results indicate that rectification can often preserve much, and in many cases all, of the desirable data in the original input.

I. Introduction

Errors and security vulnerabilities in software often occur in infrequently executed program paths triggered by atypical inputs. A standard way to ameliorate this problem is to use an anomaly detector that filters out such atypical inputs. The goal is to ensure that the program is only presented with typical inputs that it is highly likely to process without errors. A drawback of this technique is that it can filter out desirable, benign, but atypical inputs along with the malicious atypical inputs, thereby denying the user access to desirable inputs.

A. Input Rectification

We propose a new technique, *automatic input rectification*. Instead of rejecting atypical inputs, the input rectifier modifies the input so that it is typical, then presents the input to the application, which then processes the input. We have three goals: a) present typical inputs (which the application is highly likely to process correctly) to the application unchanged, b) render any malicious inputs harmless by eliminating any atypical input features that may trigger errors or security vulnerabilities, while c) preserving most, if not all, of the desirable behavior for atypical benign inputs. A key empirical observation that motivates our technique is the following:

Production software is usually tested on a large number of inputs. Standard testing processes ensure that the software performs acceptably on such inputs. We refer to such inputs as *typical inputs* and the space of such typical inputs as the *comfort zone* [33] of the application. On the other hand, inputs designed to exploit security vulnerabilities (i.e., *malicious inputs*) often lie outside the comfort zone. If the rectifier is able to automatically detect inputs that lie outside the comfort zone and map these inputs to corresponding *meaningfully close inputs* within the comfort zone, then it is possible to a) prevent attackers from exploiting the vulnerability in the software, while at the same time b) preserving desirable data in atypical inputs (either benign or malicious) for the user.

We present SOAP (Sanitization Of Anomalous inPuts), an automatic input rectification system designed to prevent *overflow vulnerabilities*. SOAP first learns a set of constraints over typical inputs that characterize a comfort zone for the application that processes those inputs. It then takes the constraints and automatically generates a rectifier that, when provided with an input, automatically produces another input that satisfies the constraints. Inputs that already satisfy the constraints are passed through unchanged; inputs that do not satisfy the constraints are modified so that they do.

B. Potential Advantages of Automatic Input Rectification

Input rectification has several potential advantages over simply rejecting malicious or atypical inputs that lie outside the comfort zone:

- **Desirable Data in Atypical Benign Inputs:** Anomaly detectors filter out atypical inputs even if they are benign. The result is that the user is completely denied access to data in atypical inputs. Rectification, on the other hand, passes the rectified input to the application for presentation to the user. Rectification may therefore deliver much or even all of the desirable data present in the original atypical input to the user.
- **Desirable Data in Malicious Inputs:** Even a malicious input may contain data that is desirable to the user. Common examples include web pages with embedded malicious content. Rectification may eliminate the exploits while preserving most desirable input from the

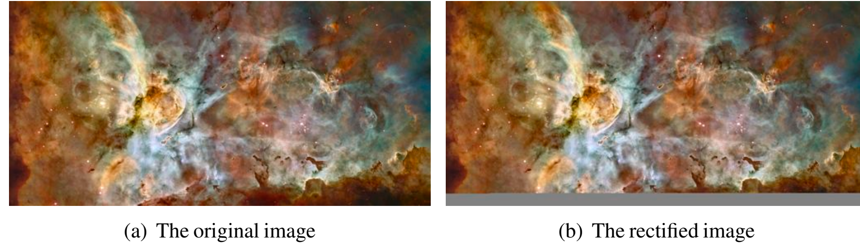

(a) The original image (b) The rectified image

Figure 1. A JPEG image truncated by the rectification.

original input. In this case the rectifier enables the user to safely access the desirable data in the malicious input.

- **Error Nullification:** Even if they are not malicious, atypical inputs may expose errors that prevent the application from processing them successfully. In this case rectification may nullify the errors so that the application can deliver most if not all of the desirable data in the atypical input to the user.

C. The Input Rectification Technique

SOAP operates on the parse tree of an input, which divides the input into a collection of potentially *nested fields*. The hierarchical structure of the parse tree reflects nesting relationships between input fields. Each field may contain an integer value, a string, or unparsed raw data bytes. SOAP infers and enforces 1) *upper bound constraints* on the values of integer fields, 2) *sign constraints* that capture whether or not an integer field must be non-negative, 3) *upper bound constraints* on the lengths of string or raw data byte fields, and 4) *correlated constraints* that capture relationships between the values of integer fields and the lengths of (potentially nested) string or raw data fields.

The taint analysis [12], [29] engine of SOAP first identifies input fields that are related to critical operations during the execution of the application (i.e., memory allocations and memory writes). The learning engine of SOAP then automatically infers constraints on these fields based on a set of training inputs. When presented with an atypical input that violates these constraints, the SOAP rectifier automatically modifies input fields iteratively until all constraints are satisfied.

D. Nested Fields in Input Files

One of the key challenges in input rectification is the need to deal with nested fields. In general, input formats may contain arbitrarily nested fields, which make inferring correlated constraints hard. Our algorithm must consider relationships between multiple fields at different levels in the tree.

Nested input fields also complicate the rectification. Changing one field may cause the file to violate constraints associated with enclosing fields. To produce a consistent rectified input, the rectifier must therefore apply a cascading sequence of modifications to correlated fields as its constraint enforcement actions propagate up or down the tree of nested fields.

E. Key Questions

We identify several key questions that are critical to the success of the input rectification technique:

- **Learning:** Is it possible to automatically learn an effective set of constraints from a set of typical benign inputs?
- **Rectification Percentage:** Given a set of learned constraints, what percentage of previously unseen benign inputs fail to satisfy the constraints and will therefore be modified by the rectifier?
- **Rectification Quality:** What is the overall quality of the outputs that the application produces when given benign inputs that SOAP has modified to enforce the constraints?
- **Security:** Does SOAP effectively protect the application against inputs that exploit errors and vulnerabilities?

We investigate these questions by applying SOAP to rectify inputs for five large software applications. The input formats of these applications include three image types (PNG, TIFF, JPEG), wave sound (WAV) and Shockwave flash video (SWF). We evaluate the effectiveness of our rectifier by performing the following experiments:

- **Benign Input Acquisition:** For each application, we acquire a set of inputs from the Internet. We run each application on each input in its set and filter out any inputs that cause the application to crash. The resulting set of inputs is the *benign inputs*. Because all of our applications are able to process all of the inputs without errors, the set of benign inputs is the same as the original set.
- **Training and Test Inputs:** We next randomly divide the collected benign inputs into two sets: the *training set* and the *test set*.
- **Potentially Malicious Inputs:** We search the CVE security database [2] and previous security papers to obtain malicious inputs designed to trigger errors and/or exploit vulnerabilities in the applications.
- **Learning:** We use the training set to automatically learn the set of constraints that characterize the comfort zone.
- **Atypical Benign Inputs:** For each application, we next compute the percentage of the benign inputs that violate at least one of the learned constraints. We call such inputs *atypical benign inputs*. In our experiments, the percentage of atypical benign inputs is less than 1.57%.

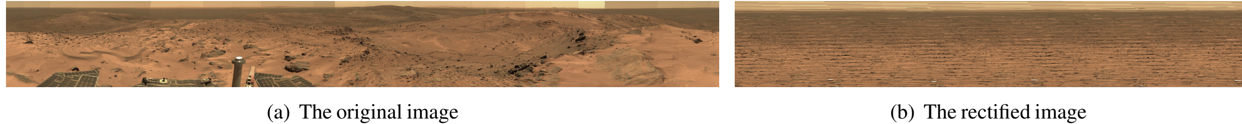

(a) The original image (b) The rectified image

Figure 2. A JPEG image twisted by the rectification

- **Quality of Rectified Atypical Inputs:** We evaluate the quality of the rectified atypical inputs by paying people on Amazon Mechanical Turk [1] to evaluate their perception of the difference between 1) the output that the application produces when given the original input and 2) the output that the application produces when given the rectified version of the original input. Specifically, we paid people to rank the difference on a scale from 0 to 3, with 0 indicating completely different outputs and 3 indicating no perceived difference. The mean scores for over 75% of the atypical inputs are greater than 2.5, indicating that Mechanical Turk workers perceive the outputs for the original and rectified inputs to be very close.
- **Security Evaluation:** We verified that the rectified versions of malicious inputs for each of these applications were processed correctly by the application.
- **Manual Code Analysis:** For each of the malicious inputs, we identified the root cause of the vulnerability that the input exploited. We examined the learned constraints and verified that if an input satisfies the constraints, then it will not be able to exploit the vulnerabilities.

F. Understanding Rectification Effects

We examined the original and rectified images or videos for all test inputs that SOAP rectified. These files are available at:

http://groups.csail.mit.edu/pac/input_rectification/

For the majority of rectified inputs (83 out of 110 inputs), the original and rectified images or videos appear identical. The mean Mechanical Turk scores for such images or videos was between 2.5 and 3.0. We attribute this to the fact that the rectifier often modifies fields (such as the name of the author of the file) that are not relevant to the core functionality of the application and therefore do not visibly change the image or video presented to the user. The application must nevertheless parse and process these fields to obtain the desirable data in the input file. Furthermore, since these fields are often viewed as tangential to the primary purpose of the application, the code that handles them may be less extensively tested and therefore more likely to contain errors.

Figures 1, 2 and 3 present examples of images that are visibly changed by rectification. For some of the rectified images (8 of 53 inputs), the rectifier truncates part of the image, leaving a strip along the bottom of the picture (see Figure 1). For the remaining inputs (19 of 110), the rectifier changes fields that control various aspects of core application functionality, for example, the alignment between pixels and the image size (see Figure 2), the image color (see Figure 3), or interactive aspects of videos. The mean Mechanical Turk scores for such images or videos vary depending on the severity of the effect. In all cases the application was able to process the rectified inputs without error to present the remaining data to the user.

G. Contributions

We make the following contributions:

- **Basic Concept:** We propose a novel technique for dealing with atypical or malicious inputs, automatic input rectification, and a prototype implementation, SOAP, which demonstrates the effectiveness of the technique.
- **Constraint Inference:** We show how to use dynamic taint analysis and a constraint inference algorithm to automatically infer safety constraints. This inference algorithm operates correctly to infer correlated constraints for hierarchically structured input files with nested fields.
- **Rectification Algorithm:** We present an input rectification algorithm that systematically enforces safety constraints on inputs while preserving as much of the benign part of the input as possible. Because it is capable of enforcing correlated constraints associated with nested input fields, this algorithm is capable of rectifying hierarchically structured input files.
- **Experimental Methodology:** We present a new experimental methodology for evaluating the significance of changes to program inputs and/or outputs. Specifically, We use Amazon Mechanical Turk [1] to evaluate the subjective perceptual quality of the outputs for rectified inputs. We present effective mechanisms to ensure the quality of the collected responses, which is a primary challenge of utilizing such crowdsourcing workforce.
- **Experimental Results:** Our results indicate that, for our set of benchmark applications and inputs, Mechanical Turk workers perceive rectified images and videos to be, in most cases, close or even identical to the original images and videos (Section V). These results are consistent with our own quantitative (Section IV) and qualitative (Section V) evaluation of the differences between the original and rectified images and videos.
- **Explanation:** We explain (Sections I-F and V) why rectification often preserves much or even all of the desirable data in rectified files.

We organize the rest of the paper as follows. Section II presents an example that illustrates how SOAP works. We describe the technical design of SOAP in Section III. We

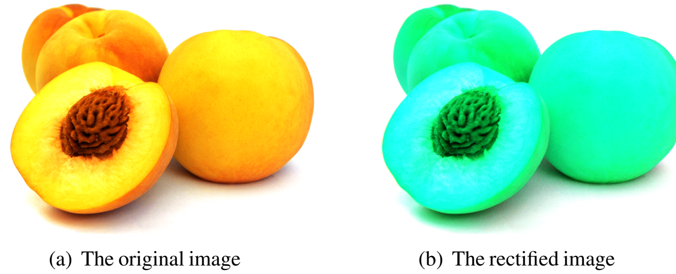

(a) The original image (b) The rectified image

Figure 3. A TIFF image whose color is changed by the rectification.

present a quantitative evaluation of SOAP in Section IV and a subjective human evaluation of SOAP in Section V. Section VI discusses related work. We conclude in Section VII.

II. Motivating Example

Figure 4 presents source code from Dillo 2.1, a lightweight open source web browser. Dillo uses libpng to process PNG files. When Dillo starts to load a PNG file, it calls the libpng callback function Png_datainfo_callback() shown in Figure 4. The function contains an integer overflow vulnerability at line 20, where the multiplication calculates the size of the image buffer allocated for future callbacks. Because png→rowbytes is proportional to the image width, arithmetic integer overflow will occur when opening a PNG image with maliciously large width and height values. This error causes Dillo to allocate a significantly smaller buffer than required. Dillo eventually writes beyond the end of the buffer.

Dillo developers are well aware of the potential for overflow errors. In fact, the code contains a check of the image size at lines 10-11 to block large images. Unfortunately, the bounds check has a similar integer overflow problem. Specific large width and height values can also cause an overflow at line 10 and thus bypass the check.

SOAP can nullify this error without prior knowledge of the vulnerability itself. To use SOAP, an application developer or system administrator first provides SOAP with a set of typical benign inputs to the application. To nullify the above Dillo error, SOAP performs following steps:

- **Understand Input Format:** SOAP provides a declarative input specification interface that enables users to specify the input format. SOAP then uses this specification to automatically generate a parser, which transforms each PNG input into a collection of potentially nested input fields. Along with the other fields of a typical PNG file, the parser will identify the locations – specifically the byte offsets – of the image width and height fields.
- **Identify Critical Fields:** SOAP monitors the execution of Dillo on training PNG inputs and determines that values in the image width and height fields influence a critical operation, the memory allocation at line 19-20. Thus SOAP marks width and height in PNG images as critical fields which may cause dangerous overflow.
- **Infer Constraints:** SOAP next infers constraints over the critical fields, including the height and width fields. Specifically, for each of these fields, SOAP infers an upper bound constraint by recording the largest value that appears in that field for all PNG training inputs.
- **Rectify Atypical Inputs:** SOAP performs the above three steps offline. Once SOAP generates constraints for the PNG format, it can be deployed to parse and rectify new PNG inputs. When SOAP encounters an atypical PNG input whose width or height fields are larger than the inferred bound, it enforces the bound by changing the field to the bound. Note that such changes may, in turn, cause other correlated constraints (such as the length of another field involved in a correlated relation with the modified field) to be violated. SOAP therefore rectifies violated constraints iteratively until all of the learned constraints are satisfied.

```
1  //Dillo's libpng callback
2  static void
3  Png_datainfo_callback(png_structp png_ptr, ...)
4  {
5    DilloPng *png;
6    ...
7    png = png_get_progressive_ptr(png_ptr);
8    ...
9    /* check max image size */
10   if (abs(png→width*png→height) >
11     IMAGE_MAX_W * IMAGE_MAX_H) {
12     ...
13     Png_error_handling(png_ptr, "Aborting...");
14     ...
15   }
16   ...
17   png→rowbytes = png_get_rowbytes(png_ptr, info_ptr);
18   ...
19   png→image_data = (uchar_t *) dMalloc(
20     png→rowbytes * png→height);
21   ...
22 }
```

Figure 4. The code snippet of Dillo libpng callback (png.c). The boldfaced code is the root cause of the overflow vulnerability.

III. Design

SOAP has four components: the *input parser*, the *execution monitor*, the *learning engine*, and the *input rectifier*. The

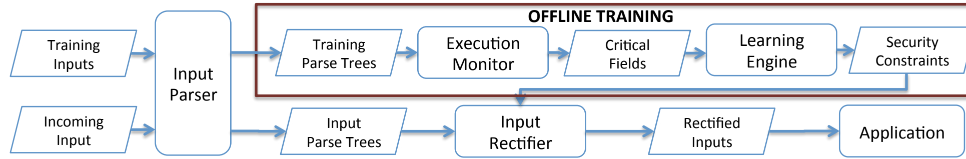

Figure 5. The architecture of the SOAP automatic input rectification system.

components work together cooperatively to enable automatic input rectification (see Figure 5). The execution monitor and the learning engine together generate safety constraints offline before the input rectifier is deployed:

- **Input parser:** The input parser ***understands input formats***. It transforms raw input files into syntactic parse trees for the remaining components to process.
- **Execution Monitor:** The execution monitor uses taint tracing to analyze the execution of an application. It ***identifies critical input fields*** that influence critical operations (i.e., memory allocations and memory writes).
- **Learning Engine:** The learning engine starts with a set of benign training inputs and the list of identified critical fields. It ***infers safety constraints*** based on the field values in these training inputs. Safety constraints define the comfort zone of the application.
- **Input Rectifier:** The input rectifier ***rectifies atypical inputs*** to enforce safety constraints. The algorithm modifies the input iteratively until it satisfies all of the constraints.

A. Input Parser

As shown in Figure 5, the input parser transforms an arbitrary input into a general syntactic parse tree that can be easily consumed by the remaining components. In the syntactic parse tree, only leaf fields are directly associated with input data. The hierarchical tree structure reflects nesting relationships between input fields. Each leaf field has a type, which can be integer, string or raw bytes, while each non-leaf field is a composite field with several child fields nested inside it. The parse tree also contains low-level specification information (e.g., endianness). The input rectifier uses this low-level information when modifying input fields.

B. Execution Monitor

The execution monitor identifies the critical input fields that should be involved in the learned constraints. Because large data fields may trigger memory buffer overflows, the execution monitor treats all variable-length data fields as critical. Integer fields present a more complicated scenario. Integer fields that influence the addresses of memory writes or the values used at memory allocation sites (e.g., calls to malloc() and calloc()) are relevant for our target set of errors. Other integer fields (for example, control bits or checksums) may not affect relevant program actions.

```
1  /header/width <= 1920
2  /header/width >= 0
3  sizebits(/text/text) <= 21112
4  /text/size * 8 == sizebits(/text/keyword)
5      + sizebits(/text/text)
```

Figure 6. A subset of constraints generated by SOAP for PNG image files.

The execution monitor uses dynamic taint analysis [12], [29] to compute the set of critical integer fields. Specifically, SOAP considers an integer field to be critical if the dynamic taint analysis indicates that the value of the field may influence the address of memory writes or values used at memory allocation sites. The execution monitor uses an automated greedy algorithm to select a subset of the training inputs for the runs that determine the critical integer fields. The goal is to select a small set of inputs that 1) minimize the execution time required to find the integer fields and 2) together cover all of the integer fields that may appear in the input files.

The execution monitor currently tracks data dependences only. This approach works well for our set of benchmark applications, eliminating 58.3%-88.7% of integer fields from consideration. It would be possible to use a taint system that tracks control dependences [10] as well.

C. Learning Engine

The learning engine works with the parse trees of the training inputs and the list of critical fields as identified by the execution monitor. It infers safety constraints over critical fields (see offline training box in Figure 5).

Safety Constraints: Overflow vulnerabilities are typically exploited by large data fields, extreme values, negative entries or inconsistencies involving multiple fields. Figure 6 presents several examples of constraints that SOAP infers for PNG image files. Specifically, SOAP infers upper bound constraints of integer fields (line 1 in Figure 6), sign constraints of integer fields (line 2), upper bound constraints of data field lengths (line 3), and correlated constraints between values and lengths of parse tree fields (lines 4-5).

These kinds of constraints enable the rectification system to eliminate extreme values in integer fields, overly long data fields, and inconsistencies between the specified and actual lengths of data fields in the input. When properly inferred and enforced, these constraints enable the rectifier to nullify the target set of errors and vulnerabilities in our benchmark.

Note that once SOAP infers a set of safety constraints for one input format, it may use these constraints to rectify inputs

for any application that reads inputs in that format. This is useful when different applications are vulnerable to the same exploit. For example, both Picasa [6] and ImageMagick [5] are vulnerable to the same overflow exploit (see Section IV). A single set of inferred constraints enables SOAP to nullify the vulnerability for both applications.

Inferring Bound Constraints: SOAP infers three kinds of bound constraints: upper bound constraints of integer fields, sign constraints of integer fields, and upper bound constraints of data field lengths. SOAP learns the maximum value of an integer field in training inputs as the upper bound of its value. SOAP learns an integer field to be non-negative if it is never negative in all training inputs. SOAP also learns the maximum length of a data field that appeared in training inputs as the upper bound of its length. SOAP infers all these constraints with a single pass over all the parse trees of the training inputs.

Inferring Correlated Constraints: SOAP infers correlated constraints in which an integer field f indicates the total length of *consecutive* children of the parent field of f. Lines 4-5 in Figure 6 present an example. The constraint states that the value of the field "$/\mathrm{text}/\mathrm{size}$" is the total length in bytes of the field "$/\mathrm{text}/\mathrm{keyword}$" and the field "$/\mathrm{text}/\mathrm{text}$", which are two consecutive nested fields inside the field "$/\mathrm{text}$".

The SOAP learning algorithm first enumerates all possible field combinations for correlated constraints initially assuming that all of these constraints are true. When processing each training input, the algorithm eliminates constraints that do not hold in the input. Our technical report [24] contains more details and peudo-code for our learning algorithm.

D. *Input Rectifier*

Given safety constraints generated by the learning engine and a new input, the input rectifier rectifies the input if it violates the safety constraints (see Figure 5). The main challenge in designing the input rectifier is enforcing safety constraints while preserving as much desirable data as possible.

Our algorithm is designed around two principles: 1) It enforces constraints only by modifying integer fields or truncating data fields—it does not change the parse tree structure of the input. 2) At each step, it attempts to minimize the value difference of the modified integer field or the amount of truncated data. It finds a single violated constraint and applies a minimum modification or truncation to enforce the constraint.

Nested input fields further complicate rectification, because changing one field may cause the file to violate correlated constraints associated with enclosing or enclosed fields at other levels. Our algorithm therefore iteratively continues the rectification process until there are no more violated constraints. In our experiments, SOAP enforces as many as 79 correlated constraints on some rectified input files.

Our algorithm has a main loop that iteratively checks the input against learned constraints. The loop exits when the input no longer violates any constraint. At each iteration, it applies a rectification action depending on the violated constraint:

- **Upper bounds of integer fields:** If the input violates the upper bound constraint of an integer field, our algorithm changes the value of the field to the learned upper bound.
- **Sign Constraints of integer fields:** If the input violates the sign constraint of an integer field, our algorithm changes the value of the field to zero.
- **Upper bounds of data field lengths:** If the input violates the upper bound constraint of the length of a data field, our algorithm truncates the data field to its length upper bound.
- **Correlated constraints:** If the value of a length indicator field is greater than the actual length of corresponding data fields, our algorithm changes the value to the actual length. If the total length of a set of data fields is longer than the length indicated by a corresponding integer field, our algorithm truncates one or more data fields to ensure that the input satisfies the constraint. Note that correlated constraints may be violated due to previous enforcements of other constraints. To avoid violating previously enforced constraints, our algorithm does not increase the value of the length indicator field or increase the field length, which may roll back previous changes.

Note that, because the absolute values of integer fields and the lengths of data fields always decrease at each iteration, this algorithm will always terminate. Our technical report [24] contains more details and pseudo-code for the algorithm.

Checksum: SOAP appropriately updates checksums after the rectification. SOAP currently relies on the input parser to identify the fields that store checksums and the method used to compute checksums. After the rectification algorithm terminates, SOAP calculates the new checksums and appropriately updates checksum fields. SOAP could use the checksum repair technique in TaintScope [37] to further automate this step.

E. *Implementation*

The SOAP learning engine and input rectifier are implemented in Python. The execution monitor is implemented in C based on Valgrind [27], a dynamic binary instrumentation framework. The input parser is implemented with Hachoir [4], a manually maintained Python library for parsing binary streams in various formats. SOAP is able to process any file format that Hachoir supports. Because SOAP implements an extensible framework, it can work with additional parser components implemented in the declarative specification interface of Hachoir to support other input formats.

IV. QUANTITATIVE EVALUATION

We next present a quantitative evaluation of SOAP using five popular media applications. Specifically, the following questions drive our evaluation:

1) Is SOAP effective in nullifying errors?
2) How much desirable data does rectification preserve?
3) How does the amount of training data affect SOAP's ability to preserve desirable data?

Application	Source	Fault	Format	Position	Related constraints
Swfdec	Buzzfuzz	X11 crash	SWF	XCreatePixMap	/rect/xmax $\leq$ 57600 /rect/ymax $\leq$ 51000
Swfdec	Buzzfuzz	overflow/crash	SWF	jpeg.c:192	/sub jpeg/.../width $\leq$ 6020 /sub jpeg/.../height $\leq$ 2351
Dillo	CVE-2009-2294	overflow/crash	PNG	png.c:142 png.c:203	/header/width $\leq$ 1920 /header/height $\leq$ 1080
ImageMagick	CVE-2009-1882	overflow/crash	JPEG,TIFF	xwindow.c:5619	/ifd[..]/img_width/value $\leq$ 14764 /ifd[..]/img_height/value $\leq$ 24576
Picasa	TaintScope	overflow/crash	JPEG,TIFF	N/A	/start_frame/content/width $\leq$ 15941 /start_frame/content/height $\leq$ 29803
VLC	CVE-2008-2430	overflow/crash	WAV	wav.c:147	/format/size $\leq$ 150

Figure 7. The six errors used in our experiments. SOAP successfully nullifies all of these errors (see Section IV-A). "Source" presents the source where we collected this vulnerability. "Fault" and "Format" present the fault type and the format of malicious inputs that can trigger this error. "Position" presents the source code file and/or line positions that are related to the root cause. "Related constraints" presents constraints generated by SOAP that nullify the vulnerability.

	Rectification Statistics							Running Time			
Input	Application	Train	Test	Field (Distinct)	Rectified	Enforced	P_{loss}	Mean	Parse	Rect.	Per field
SWF	Swfdec	3620	3620	5550.2 (98.17)	57 (1.57%)	(8.61,86)	N/A	531ms	443ms	88ms	0.096ms
PNG	Dillo	1496	1497	306.8 (32.3)	0 (0%)	(0,0)	0%	23ms	19ms	4ms	0.075ms
JPEG	IMK, Picasa	3025	3024	298.2 (75.5)	42 (1.39%)	(2.55,8)	0.08%	24ms	21ms	3ms	0.080ms
TIFF	IMK, Picasa	870	872	333.5 (84.5)	11 (1.26%)	(1.36,2)	0.50%	31ms	26ms	5ms	0.093ms
WAV	VLC	5488	5488	17.1 (16.8)	11 (0.20%)	(1.09,2)	0%	1.5ms	1.3ms	0.2ms	0.088ms

Figure 8. Benchmarks and numerical results for our experiments. The "Input" column presents the input file format. The "Application" column presents the application name (here "IMK" is an abbreviation of ImageMagick). The "Train" column presents the number of training inputs. The "Test" column presents the number of test inputs. The "Field (Distinct)" column has entries of the form X(Y), where X is the mean number of fields in each test input of each format and Y is the mean number of semantically distinct fields (i.e., fields that have different names) in each test input. The "Rectified" column has entries of the form X(Y), where X is the number of test inputs that the rectifier modified and Y is the corresponding percentage of modified test inputs. The "Enforced" column has entries of the form (X,Y), where X is the mean number of constraints that SOAP enforced for each rectified test input and Y is the maximum number of constraints that SOAP enforced over all rectified test inputs of that format. The "P_{loss}" column presents the mean data loss percentage over all test inputs of each format (see Section IV-B). The "Mean" column presents the mean running time for each test input including both parsing and rectification. The "Parse" column presents the mean parsing time for each input. The "Rect." column presents the mean rectification time for each input. The "Per field" column presents the mean running time divided by the number of input fields.

Applications and Errors: We use SOAP to rectify inputs for five applications: Swfdec 0.5.5 (a shockwave player) [7], Dillo 2.1 (a browser) [3], ImageMagick 6.5.2-8 (an image processing toolbox) [5], Google Picasa 3.5 (a photo managing application) [6], and VLC 0.8.6h (a media player) [8].

Figure 7 presents a description of each error in each application. These applications consume inputs that (if crafted) may cause the applications to incorrectly allocate memory or perform an invalid memory access. The input formats for these errors are the SWF Shockwave Flash format; the PNG, JPEG, and TIF image formats; and the WAV sound format.

Malicious inputs: We obtained six malicious input files from the CVE database [2], the Buzzfuzz project [19] and the TaintScope project [37]. Each input targets a distinct error (see Figure 7) in at least one of the examined applications.

Benign inputs: We implemented a web crawler to collect input files for each format (see Figure 8 for the number of collected inputs for each input format). Our web crawler uses Google's search interface to acquire a list of pages that contain at least one link to a file of a specified format (e.g., SWF, JPEG, or WAV). The crawler then downloads each file linked within each page. We verified that all of these inputs are benign, i.e., the corresponding applications successfully process these inputs. For each format, we randomly partitioned these inputs into two sets, the training set and the test set (see Figure 8).

A. *Nullifying Vulnerabilities*

We next evaluate the effectiveness of SOAP in nullifying six vulnerabilities in our benchmark (see Figure 7). We applied the rectifier to the obtained malicious inputs. The rectifier detected that all of these inputs violated at least one constraint. It enforced all constraints to produce six corresponding rectified inputs. We verified that the applications processed the rectified inputs without error and that none of the rectified inputs exploited the vulnerabilities. We next discuss the interactions between the inputs and the root cause of each vulnerability.

Flash video: The root cause of the X11 crash error in Swfdec is a failure to check for large Swfdec viewing window sizes as specified in the input file. If this window size is very large, the X11 library will allocate an extremely large buffer for the window and Swfdec will eventually crash. SOAP nullifies this error by enforcing the constraints $/\mathrm{rect}/\mathrm{xmax} \leq 57600$ and $/\mathrm{rect}/\mathrm{ymax} \leq 51000$, which limit the window to a size that Swfdec can handle. In this way, SOAP ensures that no rectified input will be able to exploit this error in Swfdec.

The integer overflow vulnerabilities in Swfdec occurs when Swfdec calculates the required size of the memory buffer for JPEG images embedded within the SWF file. If the SWF input file contains a JPEG image with abnormally large specified width and height values, this calculation will overflow and Swfdec will allocate a buffer significantly smaller than the required size. When SOAP enforces the learned constraints, it nullifies the error by limiting the size of the embedded image. No rectified input will be able to exploit this error.

Image: Errors in Dillo, ImageMagick and Picasa have similar root causes. A large PNG image with crafted width and height can exploit the integer overflow vulnerability in Dillo (see Section II). The same malicious JPEG and TIFF images can exploit vulnerabilities in both ImageMagick (running on Linux) and Picasa Photo Viewer (running on Windows). ImageMagick does not check the size of images when allocating an image buffer for display at magick/xwindow.c:5619 in function XMakeImage(). Picasa Photo Viewer also mishandles large image files [37]. By enforcing the safety constraints, SOAP limits the size of input images and nullifies these vulnerabilities (across applications and operating systems).

Sound: VLC has an overflow vulnerability when processing the format chunk of a WAV file. The integer field $/\mathrm{format}/\mathrm{size}$ specifies the size of the format chunk (which is less than 150 in typical WAV files). VLC allocates a memory buffer to hold the format chunk with the size of the buffer equal to the value of the field $/\mathrm{format}/\mathrm{size}$ plus two. A malicious input with a large value (such as 0xfffffffe) in this field can exploit this overflow vulnerability. By enforcing the constraint $/\mathrm{format}/\mathrm{size} \leq 150$, SOAP limits the size of the format chunk in WAV files and nullifies this vulnerability.

These results indicate that SOAP effectively nullifies all six vulnerabilities. Our code inspection proves that the learned constraints nullify the root causes of all of the vulnerabilities so that no input, after rectification, can exploit the vulnerabilities.

B. Data Loss

We next compute a quantitative measure of the rectification effect on data loss. For each format, we first apply the rectifier to the test inputs. We report the mean data loss percentage of all test inputs for each format. We use the following formula to compute the data loss percentage of a rectified input i:

$$P_{loss}^{i} = \frac{D_{loss}^{i}}{D_{tot}^{i}}$$

D_{tot}^{i} measures the amount of desirable data before rectification and D_{loss}^{i} measures the amount of desirable data lost in the rectification process. For JPEG, TIFF and PNG files, D_{tot}^{i} is the number of pixels in the image and D_{loss}^{i} is the number of pixels that change after rectification. For WAV files, D_{tot}^{i} is the number of frames in the sound file and D_{loss}^{i} is the number of frames that change after rectification. Because SWF files typically contain interactive content such as animations and dynamic objects that respond to user inputs, we did not attempt to develop a corresponding metric. We instead rely solely on our human evaluation in Section V for SWF files.

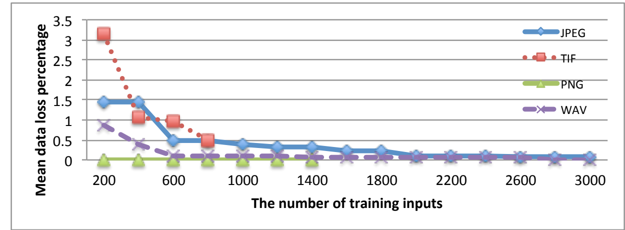

Figure 9. The mean data loss percentage curves under different sizes of training input sets for JPEG, TIFF, WAV and PNG (see Section IV-C). X-axis indicates the size of training input sets. Y-axis indicates the mean data loss percentage.

Result Interpretation: Figure 8 presents rectification results from the test inputs of each input format. First, note that more than 98% of the test inputs satisfy all constraints and are therefore left unchanged by the rectifier. Note also that both PNG and WAV have zero desirable data loss — PNG because the rectifier did not modify any test inputs, WAV because the modifications did not affect the desirable data. For JPEG and TIFF, the mean desirable data loss is less than 0.5%.

One of the reasons that the desirable data loss numbers are so small is that rectifications often change fields (such as the name of the author of the data file or the software package that created the data file) that do not affect the output presented to the user. The application must nevertheless parse and process these fields to obtain the desirable data in the input file.

C. Size of Training Input Set

We next investigate how the size of the training input set affects the rectification result. Intuitively, we expect that using fewer training inputs will produce more restrictive constraints which, in turn, will cause more data loss in the rectification. For each format, we incrementally increase the size of the training input set and record the data loss percentage on the test inputs. At each step, we increase the number of training inputs by 200. Figure 9 presents curves which plot the mean data loss percentage of the different input formats as a function of the size of the training input set.

As expected, the curves initially drop rapidly, then approach a limit as the training set sizes become large. Note that the PNG and WAV curves converge more rapidly than the TIFF and JPEG curves. We attribute this to the fact that the PNG and WAV formats are simpler than the TIFF and JPEG formats (see Figure 8 for the number of semantically distinct fields).

D. Overhead

We next evaluate the rectification overhead that SOAP introduces. Figure 8 presents the mean running time of the SOAP rectifier for processing the test inputs of each file format. All times are measured on an Intel 3.33GHz 6-core machine with SOAP running on only one core.

The results show that the majority of the execution time is incurred in the Hachoir parsing library, with the execution time per field roughly constant across the input formats (so SWF files take longer to parse because they have more fields than other kinds of files). We expect that users will find these rectification overheads negligible during interactive use.

V. Mechanical Turk-Based Evaluation

Amazon Mechanical Turk [1] is a Web-based crowdsourcing labor market. Requesters post Human Intelligence Tasks (HITs); workers solve those HITs in return for a small payment. We organized the experiment as follows:

- **Input Files:** We collected all of the TIFF, JPEG, and SWF test input files that the rectifier modified (we exclude PNG and WAV files because the original and rectified files have no differences that are visible to the user).
- **HIT Organization:** Together, the TIFF and JPEG files comprise the image files. The SWF files comprise a separate pool of video files. We partition the image files into groups, with four files per group. There is one HIT for each group; the HIT presents the original and rectified versions of the files in the group to the worker for rating. The HIT also contains a control pair. With probability 0.5 the control pair consists of identical images; with probability 0.5 the control pair consists of two completely different images. We similarly create HITs for the videos.
- **HIT Copies:** We post 100 copies of each HIT on Mechanical Turk. Each Mechanical Turk worker rates each pair in the HIT on a scale from 0 to 3. A score of 3 indicates no visible difference between the images (or videos), 2 indicates only minor visible differences, 1 indicates a substantial visible difference, and 0 indicates that the two images (or videos) appear completely different.

As with all marketplaces that involve the exchange of currency, Amazon's Mechanical Turk contains misbehaving users. For example, some workers attempt to game the system by using automated bots to perform the HITs or simply by providing arbitrary answers to HITs without attempting to perform the evaluation [22]. We used three mechanisms to recognize and discard results from such workers:

- **Approval Rating:** Amazon rates each Mechanical Turk worker and provides this information to HIT requestors. This rating indicates what percentage of that worker's previously performed HITs were accepted by other requestors as valid. We required that prospective Mechanical Turk workers have an acceptance rate of at least 95%. Using an approval rating filter provides an initial quality filter but cannot guarantee future worker performance.
- **Control Pairs:** Each HIT contains five pairs, one of which was a control pair. Half of the control pairs contained identical images or videos, while the other half contained completely different images or videos (one of the images or videos was simply null). If a worker did not correctly evaluate the control pair, we discarded the results from that worker. This technique can effectively detect bots and misbehaving workers but requires control pairs that are difficult to misinterpret.
- **Descriptions:** For each HIT, we require workers to provide a short description of the perceived differences (if any) between image or video pairs. By forcing users to provide a textual description, we help users transition from performing motor control actions (e.g., clicking on images) to cognitive executive functions. This technique helps improve the performance of legitimate workers and enables the detection of misbehaving users by monitoring for empty or nonsensical descriptions.

Format	Undetectable	Minor	Substantial	Complete
SWF	43 (1.19%)	7 (0.19%)	7 (0.19%)	0
JPEG	37 (1.22%)	3 (0.10%)	1 (0.03%)	1 (0.03%)
TIFF	3 (0.34%)	5 (0.57%)	2 (0.23%)	1 (0.11%)

Figure 10. Results of Mechanical Turk experiment. "Undetectable", "Minor" , "Substantial" and "Complete" correspond, respectively, to rectified files whose mean scores are in $[2.5, 3]$, $[1.5, 2.5)$, $[0.5, 1.5)$ and $[0, 0.5)$.

Whenever we discarded a result from a worker , we reposted a copy of the HIT to ensure that we obtained results for all 100 copies of each HIT.

Results: For each HIT h, we computed the mean scores over all the scores given by the workers assigned to h. We then used the mean scores to classify the files in h into four categories: undetectable difference (score in $[2.5, 3]$), minor difference (score in $[1.5, 2.5)$), substantial difference (score in $[0.5, 1.5)$), and complete difference (score in $[0, 0.5)$).

Figure 10 presents, for each combination of input file format and classification, an entry of the form X(Y), where X is the number of files in that classification and Y is the corresponding percentage out of all test inputs. Note that, out of 110 rectified inputs, only two exhibit a complete difference after rectification. Only 12 exhibit more than a minor difference.

To compare the Mechanical Turk results with the quantitative data loss percentage results on image files (see Section IV-B), we compute the correlation coefficient between these two sets of data. The correlation coefficient is -0.84, which indicates that they are significantly correlated ($p < 0.01$). For complex rectification effects, we find that Mechanical Turk workers can provide a more intuitive evaluation than the than quantitative data loss percentage provides. For example, only the image color in Figure 3 changes (Mechanical Turk score 1.42), but the quantitative data loss percentage reports simply that all pixels change.

Causes of Rectification Effects: When we compare the original and rectified JPEG files, we observe essentially three outcomes: 1) The rectification changes fields that do not affect the image presented to the user — the original and rectified images appear identical (37 out of 42 inputs with Turk scores in $[2.5, 3.0]$). 2) The rectification truncates part of the picture,

removing a strip along the bottom of the picture (3 out of 42 inputs with Turk scores in $[2.0, 2.3]$, see Figure 1). 3) The rectification changes the metadata fields of the picture, the pixels wrap around, and the rectified image may have similar colors as the original but with the detail destroyed by the pixel wrap (2 out of 42 inputs with Turk scores in $[0, 1)$, see Figure 2).

For TIFF files, we observed essentially four outcomes: 1) The rectification changes fields that do not affect the image presented to the user — the original and rectified images appear identical (3 out of 11 inputs with Turk scores in $[2.5, 3.0]$). 2) The rectification truncates the image, removing a strip along the bottom of the picture (5 out of 11 inputs with Turk scores in $[1.0, 2.5]$). 3) The rectification changes the color palette fields so that only the image color changes (2 out of 11 inputs with Turk scores in $[1.5, 2.0]$, see Figure 3). 4) The rectification changes metadata fields and all data is lost (1 out of 11 inputs with Turk score 0.2).

For SWF files, we observed essentially three outcomes: 1) The rectification changes fields that do not affect the video (43 out of 57 inputs with Turk scores in $[2.5, 3.0]$). 2) The rectification changes fields that only affect a single visual object in the flash video such as an embedded image or the background sound, leaving the SWF functionality largely or partially intact (3 out of 57 inputs with Turk scores in $[1.5, 2.5]$). 3) The rectification changes fields that affect the program logic of the flash video so that the rectified flash fails to respond to interactive events from users (11 out of 57 inputs with Turk scores in $[0.5, 2.6]$ depending on how important the affected events are to the users).

VI. Related Work

Input Rectification: Applying input rectification to improve software reliability and availability was first introduced by Rinard [33], who presented the implementation of a manually crafted input rectifier for the Pine email client. SOAP improves upon the basic concept by automating the fundamental components of the approach: learning and rectification.

Data Diversity: Ammann and Knight [9] propose to improve software reliability using data diversity. Given an input that triggers an error, the goal is to retry with a reexpressed input that avoids the error but generates an equivalent result. Input rectification, in contrast, may change the input (and therefore change the output). The freedom to change the input semantics enables input rectification to nullify a broader class of errors in a broader class of applications (specifically, applications for which equivalent inputs may not be available).

Anomaly Detection: Anomaly detection research has produced a variety of techniques for detecting malicious inputs [34], [23], [35], [26], [20], [30], [36]. Web-based anomaly detection [34], [23] uses input features (e.g. request length and character distributions) from attack-free HTTP traffic to model normal behaviors. HTTP requests that contain features that violate the model are flagged as anomalous and dropped. Similarly, Valeur et al [35] propose a learning-based approach for detecting SQL-injection attacks. Wang et al [36] propose a technique that detects network-based intrusions by examining the character distribution in payloads. Perdisci et al [30] propose a clustering-based anomaly detection technique that learns features from malicious traces (as opposed to benign traces). SOAP differs from anomaly detection techniques in its aim to rectify inputs and to preserve desirable data in inputs, while anomaly detection techniques simply recognize and drop potentially malicious inputs.

Signature Generation: Vigilante [14], Bouncer [13], PacketVaccine [38], VSEF [28], and ShieldGen [15] generate vulnerability signatures from known exploits. SOAP differs from such systems in its ability to nullify unknown vulnerabilities and to enable users to access desirable data in potentially malicious inputs (rather than discarding such inputs).

Critical Input, Code, and Data Inference: Snap [11] can automatically learn which input fields, code, and program data are critical to the acceptability of the output that a given application produces. Other fields, code, and data can sustain significant perturbations without changing the acceptability of the output. SOAP could use this criticality information to minimize or even eliminate changes to critical input fields in the rectification process.

Directed Fuzzing: SOAP uses taint analysis to track input fields that may trigger overflow errors. BuzzFuzz also uses taint tracing to track disparate input bytes that simultaneously reach security critical operations [19]. BuzzFuzz uses this information to perform directed fuzzing on inputs that have complex structures. Like BuzzFuzz, SOAP learns which bytes reach security critical operations. Unlike BuzzFuzz, SOAP also learns and enforces safety constraints over these bytes.

Automatic Patch: Like SOAP, ClearView [31] enforces learned invariants to eliminate errors and vulnerabilities. Specifically, ClearView learns invariants over registers and memory locations, detects critical invariants that are violated when an adversary attempts to exploit a security vulnerability, then generates and installs patches that eliminate the vulnerability by modifying the program state to enforce the invariants.

Rectification Algorithm: The SOAP rectification algorithm is inspired by automated data structure repair [17], [21], [18], which iteratively modifies a data structure to enforce data consistency defined in an abstract model. It is also possible to use data structure repair to enforce learned data structure consistency properties [16].

Evaluation with Mechanical Turk: By enabling a large-scale, low-cost human workforce, Mechanical Turk has become a viable option for a variety of experimental tasks such as training data annotation [25], computation result evaluation [22], and behavior research [32].

VII. Conclusion

Our results indicate that input rectification can effectively nullify errors in applications while preserving much, and in many cases, all, of the desirable data in complex input files.

References

[1] Amazon mechanical turk. https://www.mturk.com/mturk/welcome.

[2] Common vulnerabilities and exposures (CVE). http://cve.mitre.org/.

[3] Dillo. http://www.dillo.org/.

[4] Hachoir. http://bitbucket.org/haypo/hachoir/wiki/Home.

[5] Imagemagick. http://www.imagemagick.org/script/index.php.

[6] Picasa. http://picasa.google.com/.

[7] Swfdec. http://swfdec.freedesktop.org/wiki/.

[8] VLC media player. http://www.videolan.org/.

[9] P. E. Ammann and J. C. Knight. Data diversity: An approach to software fault tolerance. *IEEE Transactions on Computers*, 37(4):418–425, 1988.

[10] M. Attariyan and J. Flinn. Automating configuration troubleshooting with dynamic information flow analysis. In *Proceedings of the 9th USENIX conference on Operating systems design and implementation*, OSDI'10. USENIX Association, 2010.

[11] M. Carbin and M. C. Rinard. Automatically identifying critical input regions and code in applications. In *Proceedings of the 19th international symposium on Software testing and analysis*, ISSTA '10, 2010.

[12] J. Clause, W. Li, and A. Orso. Dytan: a generic dynamic taint analysis framework. In *Proceedings of the 2007 international symposium on Software testing and analysis*, ISSTA '07. ACM, 2007.

[13] M. Costa, M. Castro, L. Zhou, L. Zhang, and M. Peinado. Bouncer: securing software by blocking bad input. In *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, SOSP '07. ACM, 2007.

[14] M. Costa, J. Crowcroft, M. Castro, A. Rowstron, L. Zhou, L. Zhang, and P. Barham. Vigilante: end-to-end containment of internet worms. In *Proceedings of the twentieth ACM symposium on Operating systems principles*, SOSP '05. ACM, 2005.

[15] W. Cui, M. Peinado, and H. J. Wang. Shieldgen: Automatic data patch generation for unknown vulnerabilities with informed probing. In *Proceedings of 2007 IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2007.

[16] B. Demsky, M. D. Ernst, P. J. Guo, S. McCamant, J. H. Perkins, and M. Rinard. Inference and enforcement of data structure consistency specifications. In *ISSTA 2006, Proceedings of the 2006 International Symposium on Software Testing and Analysis*, pages 233–243, Portland, ME, USA, July 18–20, 2006.

[17] B. Demsky and M. Rinard. Data structure repair using goal-directed reasoning. In *Proceedings of the 27th international conference on Software engineering*, ICSE '05. ACM, 2005.

[18] B. Elkarablieh and S. Khurshid. Juzi: a tool for repairing complex data structures. In *Proceedings of the 30th international conference on Software engineering*, ICSE '08. ACM, 2008.

[19] V. Ganesh, T. Leek, and M. Rinard. Taint-based directed whitebox fuzzing. In *ICSE '09: Proceedings of the 31st International Conference on Software Engineering*. IEEE Computer Society, 2009.

[20] D. Gao, M. K. Reiter, and D. Song. On gray-box program tracking for anomaly detection. In *Proceedings of the 13th conference on USENIX Security Symposium - Volume 13*, SSYM'04. USENIX Association, 2004.

[21] I. Hussain and C. Csallner. Dynamic symbolic data structure repair. In *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 2*, ICSE '10. ACM, 2010.

[22] A. Kittur, E. H. Chi, and B. Suh. Crowdsourcing user studies with Mechanical Turk. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, CHI '08. ACM, 2008.

[23] C. Kruegel and G. Vigna. Anomaly detection of web-based attacks. In *Proceedings of the 10th ACM conference on Computer and communications security*, CCS '03. ACM, 2003.

[24] F. Long, V. Ganesh, M. Carbin, S. Sidiroglou, and M. Rinard. Automatic input rectification. MIT-CSAIL-TR-2011-044.

[25] M. Marge, S. Banerjee, and A. Rudnicky. Using the amazon mechanical turk for transcription of spoken language. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, 2010.

[26] D. Mutz, F. Valeur, C. Kruegel, and G. Vigna. Anomalous system call detection. *ACM Transactions on Information and System Security*, 9, 2006.

[27] N. Nethercote and J. Seward. Valgrind: a framework for heavyweight dynamic binary instrumentation. In *Proceedings of the 2007 ACM SIGPLAN conference on Programming language design and implementation*, PLDI '07. ACM, 2007.

[28] J. Newsome, D. Brumley, and D. X. Song. Vulnerability-specific execution filtering for exploit prevention on commodity software. In *NDSS*, 2006.

[29] J. Newsome and D. Song. Dynamic taint analysis: Automatic detection, analysis, and signature generation of exploit attacks on commodity software. In *Proceedings of the Network and Distributed Systems Security Symposium*, 2005.

[30] R. Perdisci, W. Lee, and N. Feamster. Behavioral clustering of http-based malware and signature generation using malicious network traces. In *Proceedings of the 7th USENIX conference on Networked systems design and implementation*, NSDI'10. USENIX Association, 2010.

[31] J. H. Perkins, S. Kim, S. Larsen, S. Amarasinghe, J. Bachrach, M. Carbin, C. Pacheco, F. Sherwood, S. Sidiroglou, G. Sullivan, W.-F. Wong, Y. Zibin, M. D. Ernst, and M. Rinard. Automatically patching errors in deployed software. In *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, SOSP '09, pages 87–102, New York, NY, USA, 2009. ACM.

[32] D. G. Rand. The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *Journal of Theoretical Biology*, 2011.

[33] M. C. Rinard. Living in the comfort zone. In *Proceedings of the 22nd annual ACM SIGPLAN conference on Object-oriented programming systems and applications*, OOPSLA '07. ACM, 2007.

[34] W. Robertson, G. Vigna, C. Kruegel, and R. A. Kemmerer. Using generalization and characterization techniques in the anomaly-based detection of web attacks. In *Proceedings of the 13 th Symposium on Network and Distributed System Security (NDSS)*, 2006.

[35] F. Valeur, D. Mutz, and G. Vigna. A learning-based approach to the detection of sql attacks. In *DIMVA 2005*, 2005.

[36] K. Wang and S. J. Stolfo. Anomalous payload-based network intrusion detection. In *RAID*, 2004.

[37] T. Wang, T. Wei, G. Gu, and W. Zou. Taintscope: A checksum-aware directed fuzzing tool for automatic software vulnerability detection. In *Proceedings of the 31st IEEE Symposium on Security & Privacy (Oakland'10)*, 2010.

[38] X. Wang, Z. Li, J. Xu, M. K. Reiter, C. Kil, and J. Y. Choi. Packet vaccine: black-box exploit detection and signature generation. In *Proceedings of the 13th ACM conference on Computer and communications security*, CCS '06. ACM, 2006.