

Virtual Static Security Analyzer for Web Applications

Mihail Brinza, Miguel Correia, João Pereira

INESC-ID, Instituto Superior Técnico, Universidade de Lisboa – Lisbon, Portugal
{mihail.brinza, miguel.p.correia, joao.d.pereira}@tecnico.ulisboa.pt

Abstract—Web applications are popular victims of injection attacks such as SQL injection and cross-site scripting. Vulnerability detection tools allow preventing these attacks but are often bound to a single language and hard to port to new languages. We propose a new approach to support the addition of new languages without much effort. In order to achieve this, our solution does not analyze the source code AST directly, instead, it traverses the source code AST and builds a generic AST (GAST) from it. Then, the tool analyzes the GAST to find vulnerabilities. This way we decouple the analysis and the source code parsing. To add support for a new language we just need to generate a parser and write a converter for that AST, which is usually less than 110 lines of code. We implemented a tool called GT with this approach. The tool currently supports four languages: Java, PHP, Python and JavaScript. It was tested against several web applications written in the same languages.

I. INTRODUCTION

In the past couple of decades, research on web security and bug finding tools have seen a big increase. This is due to the fact that software vulnerabilities can have devastating effects on companies and/or its clients [1]. In 2017, hackers have compromised the sensitive information of 145 million American customers from Equifax, one of the three major consumer credit reporting agencies in the U.S.A., leading to hundreds of millions of dollars of loss to the company [2].

Web applications are popular victims of security attacks since they accept potentially malicious user input and incorporate it into dynamically generated code. For example, a user may have to fill a form, post a comment, or submit a username and a password for authentication. The application then takes this user-provided input and inserts it into a dynamically generated program in another language (e.g., a new client-side script, or an SQL or JavaScript query to a back-end database). If the user input reaches these scripts/queries without first being validated and sanitized, then there is probably a vulnerability.

Code injection attacks such as SQL injection or cross-site scripting were considered the top security problems in 2017 by OWASP [3]. These attacks occur when a malicious user manages to inject his code into dynamically generated scripts/queries, usually by adding meta-characters to the input. By doing this, an attacker could change the behavior of the application, steal data, compromise database integrity and/or bypass authentication and access control, violating system correctness, security, and privacy properties.

Injection vulnerabilities are caused mainly by poor user input sanitization, the use of languages where it is easy to write insecure code (e.g., PHP, C) and programmers that do not have much knowledge about software security [4].

In order to detect and prevent injection attacks, we need automatic detection mechanisms. While researchers have tried many approaches over the past decades, (e.g., static and dynamic taint analysis, symbolic execution, etc.) the dominant trend is towards increasingly complex tools. However, the more complex a tool is, the worse it scales, the harder it is to maintain and understand, and the more assumptions it makes, limiting the programs it can analyze.

One big problem of the increasing complexity in vulnerability detection tools is that most of the times they are not *portable*. We use the term *portable* to refer to approaches that either support different languages or to which it is easy to add support for a new one. Although most of the programming languages we use nowadays to build web applications have a lot of similarities between them, vulnerability detection tools still seem to struggle when it comes to supporting more than one language. Many of them are wedded to a specific language [5]–[12], a specific compiled code [13]–[15] (e.g., x86 binary, bytecode), or depend on modified runtime engines [5], [16]. Porting one of these tools to another language basically requires to implement it again from scratch.

Since there is a wide range of languages available to build web applications, the lack of portability of the detection tools can be considered a problem and a limitation. However, there are some approaches that solve this issue to some degree which we discuss in Section V. Also, the tools that are specific to a language have the advantage of being able of taking into account every feature of that specific language, which in theory could lead to a more precise analysis.

In this paper we present a *new static taint analysis approach*, alongside a tool we implemented with this approach, called *Generic Taint analyzer* (GT), that aims to solve the problem of *portability* while being context-sensitive and keeping low rates of false positives and negatives. Our solution is not bound to a specific language and can be extended to support a new language with a relatively small amount of work and lines of code. Traditional static taint analyzers parse the code and then analyze the resulting abstract syntax tree (AST). The nodes that the AST consists of are specific to the parsed language, making the tools bound to that language. However, most of the languages we use nowadays to build web applications are similar. Languages usually consist of classes, attributes, functions or methods, statements, expressions, etc. Even between languages such as PHP and Java, that are apparently very different, we find many of these structural similarities. Based on this fact, our approach converts the source code AST to a simple, generic abstract syntax tree – GAST – that can represent a large set of

languages used in web applications. More importantly, GAST does not represent all the details of a language. Instead, it only has the aspects that are actually used in the analysis. This way, we use the same code to find vulnerabilities, regardless of the language being analyzed. The GAST allows us to keep the source code parsing and the analysis decoupled.

Due to the use of GAST, the analysis is divided into three steps, each one done by a specific module (*parser*, *converter*, *analyzer*):

- 1) Use a *parser* specific to the language to get the source code AST
- 2) Convert the source code AST to a GAST using a specific *converter*
- 3) Run an *analyzer* on the GAST to find vulnerabilities. In our case, the *analyzer* is a static taint analyzer

The only modules that are bound to each language are the *parser* and the *converter*; the *analyzer* is language-independent. Since parsing several different languages is a complex problem, we use ANTLR4 (ANother Tool for Language Recognition) to generate the *parser*, parse the source code and build the AST. ANTLR4 is a parser generator with a big community that provides grammars under the MIT license for virtually any language. This way, when adding support for a new language we only need to generate a new *parser* with ANTLR4 and program a new *converter*, which in our opinion requires a small amount of work.

When implementing the GT tool, we started by adding support for PHP and Java. Later, to assess portability we extended the tool to support JavaScript and Python. The converters for the four languages supported at the current stage – Java, PHP, Python and JavaScript– are less than 110 lines of code each.

The main contributions of this paper are: (1) an approach for improving portability of static security analyzers by converting the source code AST to a GAST; (2) a open-source tool that implements this approach written in Java for applications written in Java, PHP, Python and JavaScript.

II. INPUT VALIDATION VULNERABILITIES

This section briefly presents a list of vulnerabilities considered in this work. We can divide them in four categories [17] - *query manipulation*, *client-side injection*, *file and path injection* and *command injection*. The main problem of all these vulnerabilities lies in the improper validation of user input. Our work focuses on this kind of vulnerability.

A. Query Manipulation

These vulnerabilities are associated with the construction of queries or filters that are executed by some other engine (e.g., a database management system). If the query is constructed with unsanitized inputs, then it is possible to modify the normal behavior.

All vulnerabilities in this category can be prevented by sanitizing user input, so it does not contain meta-characters that can alter the behavior of the engine:

SQL Injection: vulnerability caused by the use of string building functions to create SQL queries. An attack consists of mixing normal characters with meta-characters. In the example of Listing 1, a malicious user can provide a username "admin' ---" causing the query to execute without the need of a password.

```
1 $name = $_GET['username'];
2 $pass = $_GET['password'];
3 $query = "SELECT * FROM users
4         WHERE name='$name' AND password='$pass' ";
5 $result = mysql_query($query);
```

Listing 1. PHP code vulnerable to SQL Injection

NoSQL Injection: Similarly to relational databases, non relational are also vulnerable to injection attacks.

XPath Injection: This vulnerability is very similar to SQL Injection, but in this case, data is injected in XML documents, which are often used to store data or configurations. To prevent this vulnerability, it is enough to check if the input contains malicious characters.

LDAP Injection: (Lightweight Directory Access Protocol) Injection is also exploited by providing meta-characters to string-building functions. LDAP Injection attacks aim to modify the structure of the filter and retrieve data from a directory.

B. Client-Side Injection

The vulnerabilities in this category allow an attacker to execute malicious code in the victim's browser. This kind of attack is not against the application itself but against the user and can be prevented by either sanitizing or encoding the input.

Cross-site scripting (XSS): There are three types of XSS attacks: reflected or non-persistent, stored or persistent and DOM-based. In this paper we only consider the first two types.

A program vulnerable to reflected XSS can have a single line, "echo \$_GET['user'];" The attack consists of convincing the user to click on a link to the web application with a malicious script which will be reflected by the *echo* instruction. (e.g., `www.a.pt?user=jscrip<_malicious code>jscrip<_`).

A stored XSS consists of two steps: first, the attacker inserts a malicious script in the server, then later, the server returns that script to one or more users.

Header injection: allows an attacker to break the HTTP response with "\n" and "\r". This allows the attacker to inject malicious code in headers or even a new HTTP response. It can be avoided by sanitizing these characters.

Email injection: Very similar to Header Injection, it has the goal to manipulate email components (e.g., sender, destination, message) by injecting the line termination character. In this case, sanitizing the input solves the problem as well.

C. File and Path Injection

This category consists of vulnerabilities where the user can control the path or the URL of the file included. Thus allowing him to inject malicious files (*Remote file inclusion*) or read arbitrary files from the server (*Directory/Path traversal*).

D. OS Command injection

This category consists of vulnerabilities where the user can control the input to functions that execute OS commands (e.g., *shell_exec* in PHP). Allowing the attacker to execute arbitrary code in the server.

III. GENERIC TAINT ANALYZER

In this section we explain in-depth our approach at solving the complex problem of developing a portable static taint analyzer. Our solution, besides being portable, also aims to be able to perform a context-aware analysis.

As stated before, static analyzers are most often bound to a single language and depend on the details of that language. However, there are a lot of similarities between programming languages, and our approach explores just that. If we take a look at the top trending languages in 2020, ranked by IEEE Spectrum [18], used to build web applications, we can divide them in two categories:

- 1) Dynamically typed – Languages that do not check or enforce type-safety at compile-time [19]. Instead, type-checking is done during runtime (e.g., Python, JavaScript, PHP and Ruby).
- 2) Statically typed – Type checking is done at compile-time or both (e.g., Java, C#, Go and Dart).

The languages from each category have many similarities between them. Take for instance JavaScript and PHP, which are very common on web applications. Both of them are object-oriented, have methods, functions, attributes, variables, expressions, etc... Even the *control flow statements* are practically the same (e.g., *if*, *switch*, *while*, *for*, *do while* etc.). Furthermore, the data flow is almost identical, using assignments.

Now, if we compare Java and C# the same is true. Moreover, even comparing languages between categories (e.g., Java and PHP) we see that many of their features overlap. The main difference is that we know the types when analyzing the source code. Therefore, we concluded that when it comes to static taint analysis, we can abstract many web programming languages.

With this in mind, our approach introduces a new way of building static analyzers, which consists of having a simplified, generic AST (GAST) that can represent the structure of the source code of a large set of languages found in web applications. The GAST, similarly to the micro-grammars approach [20], does not depend on every detail of the languages. Instead, it only represents what is absolutely needed to perform a data flow analysis (in our case, static taint analysis). Section III-B explains the abstractions that were made to build the GAST.

With the addition of the GAST, our analysis gains an extra step comparing to traditional tools. Usually, analyzing a program consists of first parsing the source code and then traversing the resulting AST to find vulnerabilities. By contrast, our approach adds a new, additional step, which consists of converting the source code AST to a GAST. Then, we perform the taint analysis on the GAST to detect vulnerabilities in the code. This way, the module that performs the taint analysis is completely independent of the language being analyzed.

Next, we describe the architecture and data flow of GT. Then, we present the structure of the GAST and how to build it. Finally, we discuss the features of our taint analysis and the compromises and choices that were made.

A. Architecture and Data Flow

In order to implement a tool that can support several languages simultaneously, we need a decoupled and modular architecture. The architecture of our tool, represented in Figure

1, consists mainly of four modules: *Parser*, *AST Converter*, *GAST Builder* and *Taint Analyzer*.

Parser: Takes as input the source code and produces a source code AST, specific to the language. This module is language-dependent, meaning that every time we add a new language, we need a new parser for that language. Since parsing a full-blown language is a complex task, we delegate it to ANTLR4, which is a widely used parser generator that uses *LL* for parsing [21]. ANTLR4 has a big community and provides grammars under the MIT license for virtually any language. This way, thanks to ANTLR4, our parsers consist simply of generated code. Furthermore, ANTLR4 also generates tree walkers, which we use to traverse the AST.

AST Converter: Traverses the source code AST using the generated tree walker and raises events to the *GAST Builder* (e.g., entering and exiting class, methods or functions declarations). These events allow the *GAST Builder* to create the GAST. This module is the only one needed to program whenever adding support for a new language.

GAST Builder: Reacts to the events from the *AST Converter* and internally builds the GAST. Section III-C shows more in-depth the process of building the GAST.

Taint Analyzer: Module that takes as input the GAST from the *GAST Builder* and a configuration file where we can specify the value of several parameters, such as:

- Sensitive sinks – functions that when called with *tainted* arguments may result in a vulnerability (e.g., *mysql_query* in PHP).
- Sanitization functions – For each sensitive function we can set a list of functions that can sanitize its inputs. A sanitization function is a function that cleans the input from malicious characters and whose return value is always *untainted* (e.g., one sanitization function for *mysql_query* is *mysqli_real_escape_string*). This mapping between sensitive and sanitization functions allows us precisely specify the correct sanitization function for each sensitive function, thus reducing false negatives in our analysis.
- Entry points – the function which is invoked by the server when the user makes a request. We also specify which arguments of the entry point function are *tainted*.

After traversing the GAST, GT produces a vulnerability report.

B. GAST

Similarly to micro-grammars [20], the GAST aims to be an intermediate structure that does not depend on every detail of any language. Instead, it abstracts most of the complexity by using generic statements. Each generic statement has a correspondence with a concrete statement of almost any programming language. In this section we describe the structure of the GAST by presenting every element that can be part of it.

Constant: Comprises *strings*, *integers*, *floats*, *null* values, etc. Basically, everything that is directly hardcoded in the source code is considered a constant. Since this element never changes, it is impossible to be *tainted*.

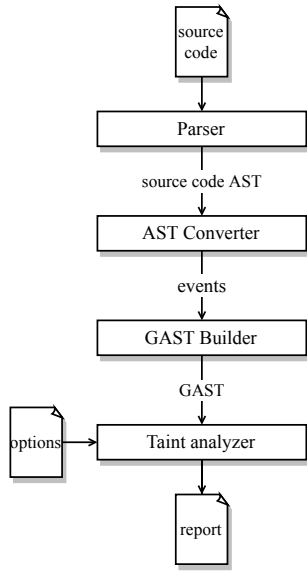


Fig. 1. The GT tool data flow and architecture

Variable: As the name says, it represents the variables in the code. Contains the name of the variable, and if the language is statically typed, it also contains the type of the variable. It can become *tainted* through assignments.

Attribute: Represents an attribute in a class. It has a name and a type. Also, it can be *tainted*.

Parameter: Represents a function or a method parameter. It can be *tainted* if it is the parameter of the entry point function, or if the argument passed to a function is *tainted*.

Expression: Represents a generic abstraction of any expression (e.g., arithmetic, logical, bitwise, comparison etc.). It consists of a list of expressions. Take for instance the Listing 2 where to each variable, *a*, *b*, *c* we assign an expression. In our representation, all these expressions are equal. They all consist of an expression containing two members: a variable named "*x*" and a constant with value "5". This allows us to abstract any operator due to the fact that in data flow analysis, operators (except the assignment operators) do not influence *taint* propagation. For instance, if the variable *x* in Listing 2 is *tainted*, then, all variables will be *tainted*, regardless of their operator.

Since an expression consists only of a list of other expressions, the *taint* propagation is a result of a logical OR of all elements in the expression. Meaning that if any element in the list is *tainted*, the whole expression is marked as *tainted*. We do not consider implicit information flows [22].

```

1 boolean a = x == 5;
2 int b = x + 5;
3 int c = x * 5;

```

Listing 2. Expression assignment examples

Function Call: This is a subtype of *expression* with the difference that it has a name, referencing the function it is invoking. The arguments are just a list of expressions that can contain anything. For instance, the function call `executeQuery(getQuery("name"))` is named *executeQuery* and has as argument another function call named *getQuery*. The latter having as argument a constant with value "*name*".

A function call is considered *tainted* if any of the following conditions is true:

- the called function is in the source code and it returns a *tainted* value
- the called function is in a library and any of its arguments is *tainted*

This element also represents method calls that do not have an object as a source (e.g., methods from the same class or any superclass that are not preceded by the keywords *this* or *super* in Java).

Assignment: Statement that represents, as the name indicates, an assignment. Consists of two expressions, one on the left-hand side and another on the right-hand side. It is the main way of propagating *taint* marks. Whenever the expression assigned on the right is evaluated as *tainted*, the mark is also propagated to the expression on the left. Usually, the expression on the left-hand side is just a variable. Listing 2 contains examples of assignment statements.

Return: Statement that represents the end of the data flow for a path. The returned value is represented by an *expression*, which can be *tainted*.

Throw: Statement very similar to the return statement. The only difference is that when a return statement is found in the callee function, the data flow is transferred to the caller function. Whilst in the case of a throw statement, the data flow is transferred to a catch block. The thrown expression can be *tainted*.

Method Call: This element represents method calls. It consists of a source, which can be an object or a class (if the method is static). Listing 3 shows examples of method calls where the source is an object and a class respectively. A method call is *tainted* if any of the following conditions is true:

- the method is not in the analyzed code and the source is *tainted*
- the method is in the source code and returns a *tainted* value

```

1 context.getUsers();
2 MyClass.myStaticMethod();

```

Listing 3. Method call examples

New: Expression that represents an object creation. It works mostly as a function call to the constructor. In most languages, the constructor has the same name as the class. But in some languages, the name of the constructor is different from the name of the class (e.g., PHP, Python). For these cases, we keep a configuration file with the names. For instance, `new Foo()` in PHP would result in the call of `__construct()` function of the class *Foo*.

Attribute access: Represents a direct access to an object attribute (e.g., `context.myProperty`).

Code block: Represents a block of code. Consists of a list of statements. Listing 2 is an example of a code block with three statements.

Conditional statement: This element represents loops (e.g., *for*, *while*, *do while* etc.). Consists of a code block and a condition, which is an *expression*.

Try Catch: Statement composed of a *try* code block, a list of *catch* code blocks and a *finally* code block.

If: Statement that represents control flow statements (e.g., *if-else*, *if-elseif-else* and *switch*). Each option in the control flow has its own code block (e.g., *if-else* has two code blocks – one for *if* and another for the *else*).

```

1  @Override
2  public void enterFunc(FuncContext ctx) {
3      gastBuilder.addFunction(ctx, ctx.name());
4  }
5
6  @Override
7  public void exitFunc(FuncContext ctx) {
8      gastBuilder.exitFunctionOrMethodDeclaration();
9  }

```

Listing 4. Function declaration example

Function: Element that represents a method or a function. It has a name, a return type, a list of parameters and a code block.

Class: Element that represents a class. It has a name, a list of attributes and a list of methods. Furthermore, it can have a superclass.

File: This is the root element of any GAST. It has a code block (for languages like PHP or JavaScript), a list of classes, a list of functions and a list of imported files (used to trace calls to imported function).

One important note to keep in mind is that not every language will use every feature from the GAST. For instance, the file element can have statements directly in the root block. Now, this is a feature that it is only used by languages such as PHP and JavaScript. By contrast, Java and C# do not allow code outside classes, so they do not make use of this feature.

C. Building the GAST

In order to build the GAST, we first parse the source code using a parser generated by ANTLR4 to obtain the source code AST. Then, we convert the AST to the GAST representation. To do this, we use a tree walker also generated by ANTLR4. The tree walker, which conceptually is a visitor, traverses every node of the AST and for each node it invokes a function when it enters or exits that node. Therefore, in order to convert the AST to GAST we need to override some of these methods in the *AST Converter* to pass state to the *GAST Builder*. This is due to the fact that the *GAST Builder* keeps the current context in a stack and these methods indicate what is the new context to push and when to pop it. This way the *AST Converter* is just a class that overrides a set of methods from the generated tree walker.

Listing 4 shows part of the PHP *AST Converter*. In the example, we override two methods from the generated tree walker. The first method is invoked when the tree walker enters a *Func* node and the second when it exits the same node. So, when the tree walker reaches a *Func* node, it invokes the method *enterFunc* from the *AST Converter*. The latter then invokes a method from the *GAST Builder* that pushes the function element to the stack. This way, the *GAST Builder* is able to keep track of the context. For instance, if the tree walker encounters a *Parameter* node while in the *Func* node, the parameter would be added to the element on the top of the stack, which in this case would be a function. The function is then popped from the stack when tree walker exits the *Func* node, invoking the method *exitFunc* which calls a method from the *GAST Builder* that pops the function (Line 8).

One important note is that we only push to the stack statements (e.g., functions, classes, assignments, conditional

statements), which are nodes in the AST, while elements such as variables, constants or parameters are not pushed. This is due to the fact that they represent leaves in the AST, so they would be pushed and popped right away. This simplification allows us to override less methods from the tree walker when writing the *AST Converter* (e.g., we do not need to override *exitVariable*).

Let us now consider Listing 5 where Line 1 corresponds to the assignment being built and Lines 3-9 to the sequence of calls made by the *AST Converter* to the *GAST Builder*. Due to lack of space, the signatures of the overridden methods are omitted. When the tree walker enters the assignment, the *AST Converter* calls a method that adds an assignment to the stack (Line 3). Next, it enters a variable and since the assignment is on the top of the stack, the variable becomes the left side of the assignment (Line 4). Then, the tree walker enters the expression "x == 5" and the *AST Converter* calls a method that adds an expression to the stack (Line 5). After that, it enters a variable and later a constant, which will both be added to the expression, since it is the top of the stack (Lines 6, 7). Finally, the tree walker first exits the expression and then the assignment, calling *exitStatementOrExpression* twice (Lines 8, 9). Figure 2 shows the GAST representation of the assignment.

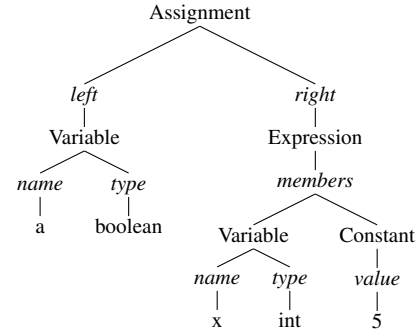


Fig. 2. GAST resulting from executing Listing 5

```

1  boolean a = x == 5;
2
3  gastBuilder.addAssignment(ctx);
4  gastBuilder.addVariable(ctx.VarName());
5  gastBuilder.addExpression(ctx);
6  gastBuilder.addVariable(ctx.VarName());
7  gastBuilder.addConstant(ctx.getText());
8  gastBuilder.exitStatementOrExpression();
9  gastBuilder.exitStatementOrExpression();

```

Listing 5. Assignment build call sequence

Converting the AST is as simple as identifying the methods needed to override from the generated tree walker, and then call the methods from the *GAST Builder*. For example, the *AST Converter* for PHP has 67 lines of code (counting only statements), and from these 67 statements, there are 29 that are different, meaning that a lot of the methods invoked are the same (e.g., *exitStatementOrExpression* is invoked 23 times).

D. Analysis features

Our approach is to use a context-aware static taint analysis to find all potential security vulnerabilities. To be able to statically find vulnerabilities, it is necessary to know what *objects* each variable may refer to, a general problem known as *pointer*,

```

1 String param = req.getParameter("name");
2 StringBuffer buf1;
3 StringBuffer buf2;
4 ...
5
6
7 buf1.append(param);
8 String query = buf2.toString();
9 con.executeQuery(query);
10
11 buf2 = buf1;
12 query = buf2.toString();
13 con.executeQuery(query);

```

Listing 6. Taint propagation example

points-to or *alias analysis* [23]. Also, the tool must perform a *path aware taint* propagation. Furthermore, it must be able to detect function/method calls between different files.

Next, we discuss the importance of the *alias analysis*. Then, we present our solution to a *path aware* analysis and the method GT uses to find function/method calls from different files. Finally, we describe our approach to handle loops.

Pointer information: To illustrate the importance of pointer information, consider the example from Listing 6. Assume that *param* is tainted and that *executeQuery* is a sensitive function. In this example, a more conservative approach may assume that *buf1* and *buf2* may reference the same object, thus marking both calls to *executeQuery* as tainted. Instead, GT traces the data flow through assignments made in each execution path. Thus, being able to identify that Line 9 is a safe call and Line 13 is a vulnerability.

Pointer analysis has been subject of much compiler research over the last decades [24], [25]. Since determining what heap object a given variable may point to is undecidable, our approach computes only an approximation based on the data flow through assignments. Meaning that, in much more complex cases, when GT is unsure to which reference an object is pointing to, it assumes that they all point to the same instance.

Path Analysis: In order to perform a precise static taint analysis it is very important to be able to track the data through different paths, ideally all of them. However, in practice, it is almost impossible since static path analysis is a very complex problem. The path is often decided at runtime due to features like *instanceof*, dynamic dispatch or reflection in Java [26]. Furthermore, these features differ between languages. Because of this, we can not have the most precise path analysis for each language. Instead, we perform an approximate path analysis based only on the *control flow statements*, ignoring their conditions. This means that we only look at the structure of the code, and for each conditional statement we propagate the data flow twice: one assuming that the flow enters that path and another assuming it does not. This assumption has the disadvantage of propagating *taint* marks through impossible paths.

Consider now the example from Listing 7 and assume that the variable *input* is *tainted*. Observing the code, we can easily identify an SQL injection vulnerability at Line 7, since at Line 3 the *query* is concatenated with the *input*. In this example, GT propagates the *taint* marks through two paths: the first executing the *if* and the second the *else*. Finally, it reports the vulnerability at Line 7. Also, it mentions that this vulnerability

```

1 String query="SELECT * FROM users WHERE name=";
2 input = null;
3 if (input != null){
4     query = query + input;
5 } else {
6     query = query + "Bob";
7 }
8 con.executeQuery(query);

```

Listing 8. Path propagation example with unreachable branch

only happens if the expression "*input != null*" at Line 2 is true. In more complex cases, for each vulnerability GT returns the call stack and the conditions that need to be met.

```

1 String query="SELECT * FROM users WHERE name=";
2 if (input != null){
3     query = query + input;
4 } else {
5     query = query + "Bob";
6 }
7 con.executeQuery(query);

```

Listing 7. Path propagation example

Let us now consider the example from Listing 8. This example is almost identical to Listing 7, with the exception of Line 2. This assignment makes the instruction from Line 4 unreachable, meaning that in practice, the code has only one possible path. However, GT has exactly the same output as the previous example: two paths and one vulnerability. This is due to the fact that GT, and static analysis tools in general, struggle with detecting whether a condition can be true or not.

This way, we perform an approximate path analysis based on the *control flow statements* and their code blocks. The advantage of this kind of path detection is that it can be applied to any language. Furthermore, from our testing, most of the times it is enough to detect vulnerabilities, even though it is not the most accurate.

Cross-file function referencing: In the last decades, web applications have become increasingly more complex, consisting of many files. For this reason, in order to perform a precise static analysis, we need to be able to perform *taint* propagation between files. However, languages have different ways of importing code. For instance, Java imports packages, which consist of a set of classes, and PHP imports files directly [27], [28].

To mitigate this problem, our approach supports two generic ways of importing code:

- 1) File inclusion – usually used by dynamically typed languages, such as PHP, Python and JavaScript. Each file has a list of imported files. When GT finds a call to a function that is not found in the file, it searches in all imported files for that function. If more than one is found, it analyses all of them.
- 2) Type tracking – works for most of the object-oriented languages, such as Java and C#. Consists of checking the type of the target of the method call and then checking if that class is in the source code. If the class is found it tries to find the method. If the method is not found, it goes to the superclass.

From our testing and analysis, most of the times GT correctly propagates the flow to other functions.

```

1 $name = $_GET["name"]
2 $query = "SELECT * FROM users WHERE name="
3 while(true) {
4     mysql_query($query);
5     $query += $name;
6 }

```

Listing 9. Vulnerability in while loop

Loops Analysis: Loops have always been tricky for static analysis tools. This is due to the fact that in many cases it is impossible to know how many times a loop will execute, if any. They are often influenced by the user. For this reason, our tool takes a simplistic approach to deal with loops, which is analyzing each loop twice. This approach helps to mitigate cases where a variable only becomes *tainted* after the first iteration. Either way, for more flexibility we left the value configurable, so we could change it depending on the program we want to analyze.

To illustrate this issue, consider Listing 9. In this example, if the *while* loop executes once, there is no vulnerability. This happens because in the first iteration *\$query* is not *tainted* upon executing *mysql_query*. However, after executing Line 5 once, *\$query* becomes tainted which makes the next call to *mysql_query* unsafe. By propagating the taint more than once, GT is able to detect this kind of vulnerability.

IV. EXPERIMENTAL EVALUATION

The objective of this section is to show that GT is capable of finding vulnerabilities in web applications written in different languages and that the effort needed to add a new language to the tool is relatively small. First, we present the results of analyzing several web applications in Java, PHP, Python and JavaScript. Then, we discuss the effort needed to add support for another language.

A. Vulnerability detection

In order to show the ability of GT to analyze and find vulnerabilities in web applications, we tested GT against two types of web applications. First, we chose 11 open source applications from GitHub that are deliberately insecure, with documented vulnerabilities. The criteria used to choose them was the number of stars that each application has on GitHub, essentially choosing the most known ones. To run the tests, we had to manually identify the entry points and sensitive functions for each application, meaning that GT analyzed each application several times, once for each entry point. Second, we also tested GT against two real-world open-source web applications. Since these applications are much bigger, we assumed that every file from the application is a web page that can receive user input, tainting the variables that might be influenced by the user (e.g., *\$_GET[*]* in PHP). Tables I and II show the results of our analysis. In the data, we only include files with the extension that we analyzed (e.g., *.java and *.php). Furthermore, we excluded comments and blank lines from the lines of code count.

GT analyzed 2015 files and 247023 lines of code and managed to find 79 documented vulnerabilities, such as SQL injection, cross-site scripting, command injection and file inclusion. The analysis times were quite low. The application that

took the longest to analyze was *PhpMyAdmin* with 82 seconds. However, the analysis times varied a lot depending on the entry point, meaning that the longer the path through which the data flows, the longer the analysis time.

In our tests, the tool had 4 false negatives in Dvna because all the dangerous code is in anonymous functions that are called by referencing variables, and our tool does not track the value of each variable during the analysis of the code. Also, we had 1 false negative when testing *PhpMyAdmin*.

GT only raised 5 false positives in the applications from Table I, due to data flow propagation through impossible paths. In total we had 12.7% of false positives and negatives.

B. Portability

Since the objective of this work is to support several languages with as little effort as possible, the portability of the tool is also a metric that we tested. To test the portability, we first implemented the tool to support PHP analysis, and then we added support for Java, which is a substantially different language. While adding support for PHP and Java we were also developing the other modules, so it is hard to tell how much time was spent strictly adding support for each language. However, later, after the tool was built, we added support for another two very popular web languages: JavaScript and Python. Table I already shows that GT is also able to find vulnerabilities in these two languages.

Table III presents the order in which the languages were added to our tool, the number of lines of each converter, the number of unique statements and the number of hours spent developing the tool to support each language. The time spent to support JavaScript and Python was roughly the same, seven hours. This is a very low number which validates our approach for building a portable vulnerability detection tool. It was possible to

This was due to the fact that the converter needed to develop to support a programming language is very simple and requires a small number of lines of code to implement. Basically, the converter consists of calls to the *GAST Builder*. In the worst case, for Java, this component has 108 lines, 30 of which are repeated (e.g., *exitStatementOrExpression()* is invoked 19 times), leaving us with 78 unique statements.

In our opinion, the main challenge when adding support for a new language is identifying which elements from the grammar are important to the analysis. After that, we just have to write the converter and some unit tests to make sure the converter works properly.

V. RELATED WORK

There is a rich body of work in the area of vulnerability detection, we just summarize the main detection methods by discussing representative papers, while leaving many others unreferences due to lack of space.

A. Taint Analysis

Taint Analysis is the most common *data flow analysis* technique used for searching for vulnerabilities [6], [8], [13]–[15], [29]–[31]. It consists of tracking the flow of sensitive

TABLE I
DELIBERATELY INSECURE WEB APPLICATIONS

Application	#loc	Language	Files	Vuln. Found	False positives	False negatives
WebGoat 8	13898	Java	320	11	1	0
Vulnado	423	Java	11	3	0	0
Dvja	950	Java	21	4	0	0
DVWA	19651	PHP	358	18	3	0
OWASP Vwa	1018	PHP	27	17	1	0
Vulnerable-node	4207	JavaScript	13	5	0	0
Dvna	771	JavaScript	14	0	0	4
Goof	571	JavaScript	8	3	0	0
NodeGoat	2697	JavaScript	49	6	0	0
Vulpy	2373	Python	57	6	0	0
Dvpwa	674	Python	21	7	0	0
Total	47233		899	78	5	4

TABLE II
REAL-WORLD WEB APPLICATIONS

Application	#loc	Lang.	Files	Vuln. found
SquirrelMail 1.5	46214	PHP	376	0
PhpMyAdmin 4.9.5	153576	PHP	740	1
Total	199790		1116	1

TABLE III
CONVERTERS SIZE AND IMPLEMENTATION EFFORT

Implementation order	Language	#loc	Unique statements	Hours to implement
1	PHP	67	49	–
2	Java	108	78	–
3	JavaScript	50	41	7
4	Python	61	49	7

information by marking user input as *tainted* and then propagate the *taint marks* recursively to the variables that are influenced by other *tainted* data. Then, it checks if *tainted* data reaches *sensitive functions* (e.g., *eval*, *my_sql_query* in PHP). If it does, then there is probably a vulnerability that could be exploited. Taint analysis can be applied to either source, binary or intermediate code.

B. Static Analysis Tools

Static analysis tools are used to automate the detection of bugs and vulnerabilities. Nowadays, they are often part of the development process, with their use being automated by continuous integration pipelines [32]. The reason is that they are a cheap way of detecting issues in code, giving the developers more confidence in their software.

SonarQube [33] is a widely used commercial static analysis tool. Performs static analysis based on a set of rules that can be defined by the user. It is able to detect bugs (e.g., possible *null references*) or bad practices in source code (e.g., empty *catch* blocks in Java). Furthermore, it also performs static taint analysis to find vulnerabilities. The taint analysis supports 4 languages, while other features support more than 20. However, since SonarQube is a commercial tool and it is not open source, we can not make any assertion about its complexity.

FlowDroid [8], is a precise static taint analyzer specifically tailored for Android and Java applications. Analyzes apps' bytecode and configuration files to find vulnerabilities. FlowDroid is precise because it models the lifecycle of android

apps and it is context, field, object and flow-sensitive. Pixy [7] performs taint analysis on PHP source code and extends it by using *alias analysis*, which takes into account the existence of aliases i.e., of two or more variable names that reference the same variable.

Notzli and Engler proposed the notion of micro-grammars [20]. They used this concept to implement an effective static analysis tool to find bugs in programs (e.g., *null pointers*, *deadlocks*). By leveraging micro-grammars, the tool is orders of magnitude less complex. This was possible by using a new source code parsing technique based on incomplete *micro-grammars*, instead of depending on every syntax detail of a language or its compiler. Traditional checking systems use parsers designed to parse a complete language syntax, thus rejecting any input that does not lead to a valid parse. On the other hand, micro-grammars parsing for bug finding has two main differences from traditional parsing: (1) When a traditional parser finds a non-matching input to its specifications, it returns an error. By contrast, when a micro-grammar parser hits a non-matching input, it simply slides forward by one token and tries again. (2) Micro-grammars allow developers to perform fine-grained input skipping by using *wildcard* non-terminals that lazily match any input up to a suffix.

The implementation of this static checker is very modular and it is composed of a lexer, parsers and checkers. There is a parser for each non-terminal, for example the parser for *C* is composed of smaller parsers for *if* statements, *while* loops, *for* loops, etc. Then all these small parsers compose the parser for *C*. This modularity brings another big advantage, the possibility of reusing a lot of these small parsers between languages, for example, C and Dart share many parsers. The use of micro-grammars and modular architecture make the tool relatively easy to port to other languages. This somehow inspired our work, since our tool also does not depend on every detail of a language.

Andromeda [34], is a demand-driven static taint analysis tool that supports Java, C# and JavaScript. It is flow and context-sensitive. Furthermore, it extends its analysis by being integrated with Framework For Frameworks (F4F), which is a solution for augmenting taint analysis with precise framework support [30].

C. Dynamic Analysis Tools

Dynamic analysis tools detect vulnerabilities at runtime. Most commonly, they rely on some kind of *code instrumentation* or transformation. Dytan [13] instruments x86 compiled code with the ability to do *taint* propagation at runtime. It is precise since it does *taint* analysis at byte level. Instrumenting x86 compiled code allows it to support any language that compiles to x86 (e.g., C and C++). Phosphor [16] applies a similar technique to Java bytecode. However, it propagates the *taint* marks at variable level, improving performance over Dytan at the cost of some precision. Moreover, it supports any language that runs on top of the Java virtual machine (e.g., Java, Kotlin or Scala).

VI. CONCLUSION

In this paper we presented an approach that aims to reduce the complexity of static security analyzers. To do this, we build a generic AST (GAST) from the source code AST. The GAST is not bound to any language and can represent a large set of languages used in web applications. This way, we perform our analysis on the GAST, allowing us to decouple the source code parsing from the vulnerability detection. The solution was implemented in a tool called GT, using Java with parsers and tree walkers generated by ANTLR4. Our evaluation shows that GT can find vulnerabilities in different types of programming languages (statically and dynamically typed) and that our goal was successfully achieved. Moreover, the evaluation also shows that our approach requires a small amount of development time to support a new programming language.

Acknowledgements We warmly thank Prof. Ibéria Medeiros for feedback on an earlier version of this work. Research supported by national funds through Fundação para a Ciência e Tecnologia (FCT) with reference PTDC/CCI-INF/29058/2017 (SEAL) and UIDB/50021/2020 (INESC-ID).

REFERENCES

- [1] R. Telang and S. Wattal, "An empirical analysis of the impact of software vulnerability announcements on firm stock price," *IEEE Transactions on Software Engineering*, vol. 33, no. 8, pp. 544–557, 2007.
- [2] H. Berghel, "Equifax and the latest round of identity theft roulette," *Computer*, vol. 50, no. 12, pp. 72–76, 2017.
- [3] J. Williams and D. Wichers, "OWASP Top 10 - 2017 rcl - the ten most critical web application security risks," OWASP Foundation, Tech. Rep., 2017.
- [4] S. Jain and M. Ingle, "Review of security metrics in software development process," *International Journal of Computer Science and Information Technologies*, vol. 2, no. 6, pp. 2627–2631, 2011.
- [5] S. Son, K. S. McKinley, and V. Shmatikov, "Diglossia: detecting code injection attacks with precision and efficiency," in *Proceedings of the 20th ACM Conference on Computer and Communications Security*, 2013, pp. 1181–1192.
- [6] I. Papagiannis, M. Migliavacca, and P. Pietzuch, "PHP Aspis: using partial taint tracking to protect against injection attacks," in *2nd USENIX Conference on Web Application Development*, vol. 13, 2011.
- [7] N. Jovanovic, C. Kruegel, and E. Kirda, "Pixy: A static analysis tool for detecting web application vulnerabilities," in *2006 IEEE Symposium on Security and Privacy (S&P'06)*, 2006, pp. 6–pp.
- [8] S. Arzt, S. Rasthofer, C. Fritz, E. Bodden, A. Bartel, J. Klein, Y. Le Traon, D. Octeau, and P. McDaniel, "Flowdroid: Precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for Android apps," *ACM SIGPLAN Notices*, vol. 49, no. 6, pp. 259–269, 2014.
- [9] P. J. C. Nunes, J. Fonseca, and M. Vieira, "phpSAFE: a security analysis tool for oop web application plugins," in *2015 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. IEEE, 2015, pp. 299–306.
- [10] G. Wassermann and Z. Su, "Static detection of cross-site scripting vulnerabilities," in *2008 ACM/IEEE 30th International Conference on Software Engineering*, 2008, pp. 171–180.
- [11] J. Dahse and T. Holz, "Simulation of built-in PHP features for precise static code analysis," in *Network and Distributed System Security Symposium*, vol. 14, 2014, pp. 23–26.
- [12] V. B. Livshits and M. S. Lam, "Finding security vulnerabilities in Java applications with static analysis," in *USENIX Security Symposium*, vol. 14, 2005.
- [13] J. Clause, W. Li, and A. Orso, "Dytan: a generic dynamic taint analysis framework," in *Proceedings of the ACM International Symposium on Software Testing and Analysis*, 2007, pp. 196–206.
- [14] J. Newsome and D. X. Song, "Dynamic taint analysis for automatic detection, analysis, and signature generation of exploits on commodity software," in *Network and Distributed System Security Symposium*, vol. 5, 2005, pp. 3–4.
- [15] M. G. Kang, S. McCamant, P. Poosankam, and D. Song, "DTA++: dynamic taint analysis with targeted control-flow propagation," in *Network and Distributed System Security Symposium*, 2011.
- [16] J. Bell and G. Kaiser, "Phosphor: Illuminating dynamic data flow in commodity jvms," *ACM Sigplan Notices*, vol. 49, no. 10, pp. 83–101, 2014.
- [17] I. V. de Sousa Medeiros, "Detection of vulnerabilities and automatic protection for web applications," Ph.D. dissertation, Doctoral dissertation, Universidade de Lisboa, 2016.
- [18] S. Cass, "The top programming languages: Our latest rankings put Python on top-again-[careers]," *IEEE Spectrum*, vol. 57, no. 8, pp. 22–22, 2020.
- [19] L. Tratt, "Dynamically typed languages," *Advances in Computers*, vol. 77, pp. 149–184, 2009.
- [20] F. Brown, A. Nötzli, and D. Engler, "How to build static checking systems using orders of magnitude less code," *ACM SIGPLAN Notices*, vol. 51, no. 4, pp. 143–157, 2016.
- [21] T. Parr, *The Definitive ANTLR 4 Reference*, 2nd ed. Pragmatic Bookshelf, 2013.
- [22] D. E. Denning, "A lattice model of secure information flow," *Communications of the ACM*, vol. 19, no. 5, pp. 236–243, 1976.
- [23] M. Sridharan, S. Chandra, J. Dolby, S. J. Fink, and E. Yahav, "Alias analysis for object-oriented programs," in *Aliasing in Object-Oriented Programming. Types, Analysis and Verification*. Springer, 2013, pp. 196–232.
- [24] J. Späth, L. Nguyen Quang Do, K. Ali, and E. Bodden, "Boomerang: Demand-driven flow-and context-sensitive pointer analysis for Java," in *30th European Conference on Object-Oriented Programming*, 2016.
- [25] M. Hind, "Pointer analysis: Haven't we solved this problem yet?" in *Proceedings of the 2001 ACM SIGPLAN-SIGSOFT workshop on Program analysis for software tools and engineering*, 2001, pp. 54–61.
- [26] C. Hammer, R. Schaade, and G. Snelting, "Static path conditions for Java," in *Proceedings of the third ACM SIGPLAN workshop on Programming languages and analysis for security*, 2008, pp. 57–66.
- [27] A. Rountev, S. Kagan, and M. Gibas, "Static and dynamic analysis of call chains in Java," in *Proceedings of the 2004 ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2004, pp. 1–11.
- [28] M. Hills, P. Klint, and J. J. Vinju, "Static, lightweight includes resolution for PHP," in *Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering*, 2014, pp. 503–514.
- [29] D. Evans and D. Larochelle, "Improving security using extensible lightweight static analysis," *IEEE Software*, vol. 19, no. 1, pp. 42–51, Jan/Feb 2002.
- [30] M. Sridharan, S. Artzi, M. Pistoia, S. Guarnieri, O. Tripp, and R. Berg, "F4F: taint analysis of framework-based web applications," in *Proceedings of the 2011 ACM International Conference on Object Oriented Programming Systems Languages and Applications*, 2011, pp. 1053–1068.
- [31] I. Medeiros, N. F. Neves, and M. Correia, "DEKANT: a static analysis tool that learns to detect web application vulnerabilities," in *Proceedings of the 25th International Symposium on Software Testing and Analysis*, 2016, pp. 1–11.
- [32] S. M. Mohammad, "Continuous integration and automation," *International Journal of Creative Research Thoughts*, pp. 2320–2882, 2016.
- [33] G. A. Campbell and P. P. Papapetrou, *SonarQube in action*. Manning Publications Co., 2013.
- [34] O. Tripp, M. Pistoia, P. Cousot, R. Cousot, and S. Guarnieri, "Andromeda: Accurate and scalable security analysis of web applications," in *International Conference on Fundamental Approaches to Software Engineering*. Springer, 2013, pp. 210–225.