



HAL
open science

An Invariant-based Approach for Detecting Attacks against Data in Web Applications

Romarc Ludinard, Eric Totel, Frédéric Tronel, Vincent Nicomette, Mohamed
Kaâniche, Eric Alata, Rim Akrouf, Yann Bachy

► **To cite this version:**

Romarc Ludinard, Eric Totel, Frédéric Tronel, Vincent Nicomette, Mohamed Kaâniche, et al.. An Invariant-based Approach for Detecting Attacks against Data in Web Applications. International journal of secure software engineering, 2014, 5 (1), pp.19-38. 10.4018/ijssse.2014010102 . hal-01083296

HAL Id: hal-01083296

<https://inria.hal.science/hal-01083296>

Submitted on 17 Nov 2014

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

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

An Invariant-based Approach for Detecting Attacks against Data in Web Applications

Romarc Ludinard¹, Éric Totel¹, Frédéric Tronel¹, Vincent Nicomette^{2,3}, Mohamed Kaâniche^{2,4},

Éric Alata^{2,3}, Rim Akrouit^{2,4}, Yann Bachy^{2,4}

¹Supélec, Avenue de la Boulaie, Cesson-Sevigné, France

²CNRS, LAAS, 7 Avenue du Colonel Roche, BP 54200, 31031 Toulouse Cedex 4, France

³Université de Toulouse, INSA, LAAS, F- 31400 Toulouse Cedex 4, France

⁴Université de Toulouse, LAAS, BP 54200, 31031 Toulouse, France

Abstract

RRABIDS (Ruby on Rails Anomaly Based Intrusion Detection System) is an application level intrusion detection system (IDS) for applications implemented with the Ruby on Rails framework. The goal of this intrusion detection system is to detect attacks against data in the context of web applications. This anomaly based IDS focuses on the modelling of the normal application profile using invariants. These invariants are discovered during a learning phase. Then, they are used to instrument the web application at source code level, so that a deviation from the normal profile can be detected at run-time. This paper illustrates on simple examples how the approach detects well-known categories of web attacks that involve a state violation of the application, such as SQL injections. Finally, an assessment phase is performed to evaluate the accuracy of the detection provided by the proposed approach.

Keywords: intrusion detection, invariant discovery, invariant checking, and web applications

An Invariant-based Approach for Detecting Attacks against Data in Web Applications

Nowadays there is an important trend to transform traditional applications into their equivalent online counterpart, making possible for end users to remotely interact with their information systems through their Web browser. Such web-based applications offer all functionalities required by the end-user (e.g., webmail, online calendars, e-business, e-banking...). This paradigm shift did not lead to a significant change of the methodology and processes employed during the application development. Hence, the design faults that are commonly encountered with traditional applications are still present in these new web-based applications. As illustrated by recent statistics (see e.g., (IBM, 2012), these applications still exhibit a high number of vulnerabilities. They have become one of the preferred targets for the attackers. The consequences of these attacks are at least as critical as they were with traditional applications. Worse, the corresponding vulnerabilities can now be remotely exploited.

Several mechanisms have been developed to protect web applications against potential attacks. These mechanisms can be used at the network level (such as applicative firewalls), and at application level (such as input sanitization techniques). However such prevention mechanisms remain largely insufficient, as they do not take into account the semantic of the applications they are supposed to protect, thus providing a rather low protection. Hence, it is important to propose mechanisms that could more efficiently detect attacks affecting the integrity of the application state.

In this paper, we propose an approach to anomaly based intrusion detection at the application level that focuses on violations of the application state. This approach is based on the automatic generation of invariants that are discovered during a learning phase and verified during the execution of the application. The paper extends the results presented in (Ludina et al., 2012) by providing more details about the implementation of the approach and the generation of the invariants. First Section **State of the Art** presents previous work about intrusion detection that can be applied in the context of web applications. Section **Context and Case Study** presents the context and a case study of our work. In Section **Invariant Constraints in an Application**, a typical example illustrates the ideas we are relying on. Then, a general overview our approach is outlined in Section **Invariant based Detection Model**, together with a discussion of how it can be applied to web applications (Section **Ruby On Rails Implementation**). The last part of the paper (Section **IDS Assessment**) shows how the intrusion detection mechanisms have been assessed, in order to demonstrate that they accurately detect the types of attacks we focus on.

State of the Art

Most work in the context of web attack detection focuses on abnormal network packets (Robertson, Vigna, Kruegel, & Kemmerer, 2006) or requests (Vigna, Robertson, Kher, & Kemmerer, 2003) and do not take into account the state of the web application itself. We believe that application level mechanisms can help improving the intrusion detection performance as they are able to take advantage of the internal state of the monitored program. Indeed, they have access to all the internal data structures and algorithms used by the program.

Three types of approaches can be distinguished for detecting intrusions at application level: the first approach focuses on the correctness of the Control Flow Graph of the program such as (Abadi, Budi, Erlingsson, & Ligatti, 2005) and (Kiriansky, Bruening, & Amarasinghe, 2002) and consists in verifying that the actions in the program are executed in a correct order.

This work does not permit to detect attacks on data. The second approach focuses on the correctness of the data manipulations during program execution such as in (Akritidis, Cadar, Raiciu, Costa, & Castro, 2008) or in (Castro, Costa, & Harris, 2006). In this case, a data-flow graph is computed prior to the execution that contains, for each data item read by an instruction, the set of instructions that may have written its current value. This data-flow graph is then used at run time to verify the integrity of the process data flow that really occurs. If a vulnerability is exploited in the monitored program that corrupts some data, a deviation from the data-flow graph will be detected the next time this data is read. A third approach, as defined by (Sarrouy, Totel, & Jouga, 2009), consists in checking the correctness of the data used by the program during its execution, rather than verifying the consistency of the data flow. This approach permits to detect a different kind of attacks on data compared to the data-flow approach (such as the use of incorrect values in legitimate variables). In (Sarrouy et al., 2009) the application normal profile is modelled as a set of constraints on the data used by the application. This model is built during a two-fold learning phase. First the data items used by the application are logged. Then, a set of invariants is generated by a dedicated tool stemming from the testing field called Daikon (Ernst et al., 2007). These invariants are then checked at run-time to perform the detection.

All these approaches are relevant in the context of both traditional and web applications. For instance, similar approaches have already been investigated by (Cova, Balzarotti, Felmetsger, & Vigna, 2007) in the context of PHP applications. In this work the authors focus both on the web application control flow and on the correctness of the data items used by the program. (Felmetsger, Cavedon, Kruegel, & Vigna, 2010) is an extension of this work to JSP. The authors propose to build the application normal behaviour during a learning phase where the invariant behaviour of the application is modelled using the Daikon tool. Then a model checking approach is used to verify if the invariants can be violated by the application, leading thus to the discovery of vulnerabilities in the application. This interesting approach suffers from several limitations, such as the discovery of invariants restricted to only inputs and outputs of the functions executed by a servlet, and the lack of consideration of temporal relationships between the variables used by the program.

In our work, we focus on the correctness of data only, but at a finer granularity (i.e., at the instruction execution level) like in (Sarrouy et al., 2009), and we study the evolution of variables with respect to time so as to discover invariants on these variables. Compared to Swaddler developed by (Cova et al., 2007), we additionally introduce the notion of tainted data (Cavallaro & Sekar, 2008) to reduce the set of variables we are checking.

In order to demonstrate the feasibility of our approach at application level, we consider the example of applications implemented with the Ruby on Rails framework. This framework is receiving increasing interest by the software development community. In the next section, we present the main concepts of Ruby on Rails that are needed to explain our intrusion detection approach, and we present the application example used as a case study in our work.

Compared to previous work, the approach proposed in this paper presents innovative aspects:

- In all the work dedicated to intrusion detection that uses Daikon, the authors use the available Daikon front-ends or extend them. Thus, they obtain invariants on inputs or outputs of the methods or functions called. In our approach, we observe the execution of all single instructions in the interpreter and we obtain invariants on all the variables manipulated by all

instructions executed by the interpreter. As a result of the previous point, much more invariants can be found compared to other approaches such as (Cova et al., 2007), (Cavallaro & Sekar, 2008) and thus our approach should obtain a better coverage for the detection of attacks.

- Our approach uses Daikon in a non-conventional way to extract from the execution traces temporal relationships on the program variables, which is a completely new functionality for a work based on Daikon.

- The implementation we propose does not imply the modification of the interpreter, compared to other approaches, providing thus an elegant implementation of invariant discovery and checking.

Context and Case Study




Figure 1. Rails MVC framework in the context of our application

Ruby on Rails (<http://rubyonrails.org/>) (RoR) is a Model-View-Controller framework that eases the way to design web applications. The choice of RoR as a case study was dictated by the fact that it provides introspection mechanisms that make it possible to observe the behaviour of the application without modifying the Ruby interpreter. A typical RoR application is composed of several controllers (like shown by Figure 1). A request to an application determines which Rails controller must be used, and which action of this controller must be called. When one of these controllers is called, it interacts with the model to access the data usually stored into a database. Using this data, it generates a view that is returned to the user as result of its request.

Since all the requests sent to the application are processed by a controller (more precisely by an action in a given controller), we chose to apply our detection model at the controller level.

In this work we have applied our approach to an e-commerce application designed by Kereval (www.kereval.com) in the context of the ANR (French National Research Agency) funded project DALI (Design and Assessment of application Level Intrusion detection systems). This application called *Insecure* intentionally exhibits the common vulnerabilities of a web application. More precisely, it is vulnerable to the following attacks: SQL injection, cross-site scripting, request parameter modification and cookie manipulation. This e-commerce application has been designed without any knowledge of the intrusion detection mechanisms that have been subsequently designed to protect it. This implies that it was not developed with the objective to emphasize the accuracy of our approach.

Invariant Constraints in an Application

This section shows on a simple example drawn from the *Insecure* application how its normal behaviour can be characterized using invariants on variables, and how an attack can violate at run-time one or more such invariants. We claim that attacks that violate the application state integrity can be detected by checking invariants at run-time.

```
class UsersController < ApplicationController
  ...
  def login
    if request.post?
      # whole user is stored in the session
      if session[:user] = User.authenticate(params[:user][:login],
                                           params[:user][:password])
        flash[:notice] = 'You have been successfully logged in.'
        if session[:user].admin
          redirect_to :controller => '/admin/home', :action => 'index'
        else
          redirect_to :controller => '/user/home', :action => 'index'
        end
      else
        flash.now[:error] = 'Login failed'
      end
    end
  end
  ...
end
```

Figure 2. Login Action in the Users Controller

Figure 2 illustrates the Ruby code used in the application to authenticate a user. This code is located in the action *login* of the controller *Users*.

```

class User < ActiveRecord::Base
  ...
  def self.authenticate(login, pass)
    user=find(:first,:conditions => "login = '#{login}' and password =
'#{pass}'")
    end
  ...
end

```

Figure 3. SQL request in the user Model

One can notice that the user's login and password are passed as parameters to the action (and stored in the *params* hashtable variable). These parameters are in turn used by a method called *User.authenticate*. This method performs a request to the database (see Figure 3) to retrieve the actual user corresponding to these parameters. The normal behaviour of this method consists in returning a user object stored in the session whose attribute login (*session[:user].login*) is equal to the login passed as parameter (*params[:user][:login]*). As a consequence, the relation *session[:user].login == params[:user][:login]* is an invariant of the application.

If an attacker performs an SQL injection by setting "'OR '1'='1'" as user login or password, he will be authenticated as the first user in the database which usually corresponds to the administrator. As a result, the previously mentioned invariant would be violated since *session[:user].login != params[:user][:login]*. This example illustrates a rather simple kind of invariants, in that sense that it links the output produced by a method call with its inputs. This type of invariants is well known and may even be discovered by static analysis tools (Karr).

```

class User::OrdersController < User::UserController
  ...
  def new
    @addresses = Address.all(:conditions => {:user_id => @user.id, :active =>
true})

    if @addresses.empty?
      flash[:warning]="No delivery address available."
      redirect_to :controller => "addresses", :action => "index"
    else
      @amount = params[:amount]
      @order = Order.new

      respond_to do |format|
        format.html # new.html.erb
        format.xml { render :xml => @product }
      end
    end
  end
  ...
end

```

Figure 4: Action called for a new order

We claim that other types of relationships can be found and should be investigated. In the approaches based on Daikon such as (Felmetsger et al., 2010) no temporal relationships can be found between successive requests. The authors even claim that such relationships do not exist unless the programmer uses global storage variables such as session variables. In fact, bad

programming practice can produce information flows from a request to another, for example by using page parameters. This is illustrated by the figures 4 and 5 that show two actions of the *Order* controller that are successively called when a client orders items from the e-commerce application. The first action is the *new* action that is in charge of presenting a summary of its order to the client. The client must validate all the needed information in the form (a valid postal address, a valid credit card number, etc.). Once the client has validated this form, the second action (*create* action) is triggered. The amount (price) of this order is transmitted as a parameter of the action. Thus an invariant can be deduced by observing the normal behaviour of the application. Namely, the variable *param[:amount]* in the *create* action must carry the same value as the variable *@amount* got in the *new* action. A possible attack against the application could consist in changing the amount parameter of the second request (using a plugin such as *TamperData* within the client navigator). However, if the previous invariant had been checked at execution, this attack would have been detected. The consideration of this particular kind of invariants is a major contribution of this paper. As far as we know RRABIDS is the first tool to propose this kind of invariants.

```

class User::OrdersController < User::UserController
  ...
  def create
    params[:order][:user_id] = @user.id
    params[:order][:card_expiration] = Date.parse(
      "#{params[:card_expiration][:year]}/#{params[:card_expiration][:month]}/1")
    @order = Order.new(params[:order])
    @amount = params[:amount]

    respond_to do |format|
      if @order.save
        save_items
        session[:cart] = nil
        flash[:notice] = "Your order has been successfully placed"
        format.html{redirect_to[:user, @order] }
        format.xml{render :xml=>@order, :status=>:created, :location=>@order }
      else
        @addresses=Address.all(:conditions=>{:user_id=>@user.id, :active =>
true})
        format.html{render :action=>"new" }
        format.xml{render :xml=>@order.errors, :status =>
:unprocessable_entity}
      end
    end
  end
  ...
end

```

Figure 5: Action called after order validation

In theory, this should be the role of the programmer to define relations between variables and check them at run-time. But in practice, very few programmers perform such checks unless they are forced to. Our work consists in automatically building such relationships and instrumenting the source code to insert such checks.

Invariant based Detection Model

Our work builds on the approach proposed by Sarrouy (Sarrouy et al., 2009) which considered C applications. It is based on the dynamic discovery of invariants on traces that are generated during the execution of the program. However, compared to (Sarrouy et al., 2009), we not only show that the behaviour of the program can be modelled as invariants, but we demonstrate that these invariants can be used at execution time to detect attacks. Three types of invariants are considered to describe: 1) relations between variables at a given execution point; 2) relations between the different values held by a given variable at different points of execution; and 3) relations between different variables at different points of execution.

Clearly, execution points play a central role in these objectives. In this paper, we define an execution point p to be a triple that consists of (i) the value of the program counter (pc), (ii) the state of the memory (mem) associated with the executed program (stack, heap, etc) and (iii) the logical time when this execution point occurs: $p=(pc, mem, t)$.

Besides, it is also necessary to define precisely the notion of a program state. For a given execution and a given execution point, the state of a program consists in the set of all variables together with their associated values that are available at this point (global or local variables). Let σ be an execution of a program. We define $S_{p,\sigma}$ to be the state of this program at execution point p during execution σ . More precisely we model $S_{p,\sigma}$ as a partial function from the set of all possible variables to their definition sets. Hence $S_{p,\sigma}(v)$ is the value taken by variable v during execution σ at point p if v is defined or \perp otherwise (i.e., if v is not defined).

Since we want to discover invariants that are valid for all possible executions, we should not limit ourselves to a single given execution (such as σ). Instead we want to consider a set of possible executions that we denote Σ . For two executions σ_1 and σ_2 and two execution points $p_1 \in \sigma_1$ and $p_2 \in \sigma_2$, we say that $p_1 = (pc_1, mem_1, t_1)$ is equivalent to $p_2 = (pc_2, mem_2, t_2)$ denoted $p_1 \equiv p_2$, if and only if their program counters are equal, that is if $pc_1 = pc_2$. Hence, two execution points are equivalent if they correspond to the same position in the program code (same program counter) without taking into account the state of their memories and their occurrence time. We introduce this equivalence relation to be consistent with the method used by Daikon to discover invariants. Indeed, Daikon analyzes traces that contain the values taken by possibly different variables at specific program counters. Hence, for Daikon two execution points are equivalent as long as they concern the same program counter. More precisely, three kinds of default program counters are defined by Daikon: (1) function entry points, (2) function exit points, and (3) any instruction within the body of a function. Note that in (Felmetsger et al.) only the first two are taken into account to compute invariants.

We finally introduce the notion of *unified* state for a given variable v over the set of considered executions Σ at a given program counter pc to be the function defined by:

$$S_{pc}(v) = \bigcup_{\sigma \in \Sigma} \bigcup_{p \in \sigma | p=(pc, *, *)} S_{p,\sigma}(v)$$

The semantic of a *unified* state is as follows: given a program counter pc , a unified state describes the values that can be associated with any accessible variable at this point of the program (that is for all equivalent execution points). From this definition we can deduce the set of all variables V_{pc} at a given execution point $V_{pc} = \{v | S_{pc}(v) \neq \perp\}$.

This set of variables contains all variables that can be used by Daikon to compute an invariant for a given program counter. Two observations can be raised about this model:

- First, it requires to save the values of all the variables each time the program counter evolves. Since this would require a very large amount of memory, we need to reduce this set to a finite and usable set of variables. This is discussed in Section Reducing the Set of Variables;

- Second, automatically finding invariants on these variables requires to define what types of invariant constraints are actually needed; this set of constraints is potentially dependent on the model of attacks we intend to detect. This is the purpose of Section Classes of Invariants.

Reducing the Set of Variables




Figure 6. Variables causally dependent on user inputs

To reduce the set of variables to be logged, we emphasize two considerations:

- First, all variables are not of interest for intrusion detection: an attacker needs to interact with the application to attack it, and thus he needs to modify its inputs. The data used by the application are altered by propagation of the modified inputs. As a consequence, the variables that are of interest are those that depend on the external inputs of the application. This set of variables can be defined formally using the notion of causal dependencies.

- Even if we restrict the set of variables to the variables that depend on external inputs, the application states to be saved are potentially still very large. Consequently, we will only keep in memory all the evolutions of the variables that occur within a given window of time. In this work, we restrict the time window to two consecutive actions performed by a single user identified by its session id (see Figure 6).

Using the notion of causal dependencies as defined in (d'Ausbourg, 1994), we say that a variable v_1 at time t_1 is causally dependent on a variable v_0 at time t_0 (noted $(v_0, t_0) \rightarrow (v_1, t_1)$) if there exists an information flow from (v_0, t_0) to (v_1, t_1) .

The set of variables that influence a variable (v_0, t_0) is called the *causality cone* and denoted $cause(v_0, t_0) = \{(v_i, t_i) \mid (v_i, t_i) \rightarrow^* (v_0, t_0)\}$ (where \rightarrow^* is the transitive closure of the relation \rightarrow).

This definition can be extended using the execution point we have previously defined: for a given variable v_0 at point $p_0 = (pc_0, mem_0, t_0)$, with the variable $v_0 \in V_{pc_0}$,

$$cause(v_0, p_0) = \{(v, p) \mid (v, p) \rightarrow^* (v_0, p_0)\}, \text{ with } (v, p) \rightarrow (v_0, p_0) \Leftrightarrow (v, t) \rightarrow (v_0, t_0)$$

We then define the notion of a *tainted* variable: informally a variable is tainted if it depends on the external inputs of the program (e.g., for a web application, user inputs, data coming from the database or external configuration files are external inputs). Formally, it means that the causality cone of this variable at a point p of the program contains an input of the program.

We can finally define the *tainted* function $T_{p,\sigma}$ of a variable v

$$T_{p,\sigma}(v) = \begin{cases} S_{p,\sigma}(v) & \text{if } v \text{ is tainted} \\ \perp & \text{otherwise} \end{cases}$$

We thus obtain the set of values that can be taken by a tainted variable at a given program counter:

$$T_{pc}(v) = \bigcup_{\sigma \in \Sigma} \bigcup_{p \in \sigma \mid p=(pc, *, *)} T_{p,\sigma}(v)$$

As previously, we can define the set of tainted variables accessible at a given program counter. We denote this set as $TV_{pc} = \{v \mid T_{pc}(v) \neq \perp\}$.

Classes of Invariants

Once the set of variable values has been built during an observation phase, we need to automatically generate the set of invariant constraints on this set of values that characterize the normal behaviour of the application. We selected a number of different invariants that could be useful to detect an intrusion inside the application:

- a variable at a specific execution point is always constant; such invariant could be useful to detect the integrity violation of some variables, e.g., some constant values stored in the database to configure the application; this type of invariant consists in finding a relation for a given program counter on all values of v in $T_{pc}(v)$.

- a variable at a specific execution point holds values that belong to small set of possible values; this type of invariants could be useful to detect the corruption of some data provided by a user, e.g., a parameter which describes a number of possible options for the execution of a request; this type of invariant consists in finding a relation for a given program counter on all values of v in $T_{pc}(v)$.

- a variable at a specific execution point is always equal to another variable at another point; this type of invariant puts in relation some variable values that could be incorrectly modified during an attack, violating the state of the application. For example, the SQL injection described in Section **Invariant Constraints in an Application** implies such state violation. This type of invariant denotes a link between two sets of variables TV_{pc1} and TV_{pc2} where $pc1$ and $pc2$ are two program counters.

- an order relationship between two variables at different execution points ($<$ or $>$); like the equality invariant, this relationship describes a normal state of a program that could be

altered by an attack that modifies the value of an application variable. This type of invariant denotes a link between two sets of variables TV_{pc1} and TV_{pc2} where pc_1 and pc_2 are two program counters.

In practice, complex invariants could be computed on the set of tainted variables, e.g., some linear relations between variables. However, the experience proved us that this restricted class of invariants is sufficient to detect the commonly known web application attacks. Moreover, the implementation actually shows that these four types of invariants generate a considerable number of actual invariants even on a small application.

The invariants that we discover must not only be true for a given execution σ , but also for all the executions of Σ . In practice, the first executions lead to the discovery of a large set of invariants. Then, this set is reduced by invalidating them on further executions of Σ .

Ruby On Rails Implementation

We introduced in Section **Context and Case Study** the Ruby on Rails framework in the context of which we built our prototype. The choice of this framework was essentially dictated by the functionalities available in the Ruby language. In order to carry out an implementation of our model, we require: (1) to be able to intercept the execution of any instruction in the interpreter and (2) to be able to access the current execution context (i.e., the variables that are part of this context) at each execution point in order to log the state of the program. The Ruby on Rails framework and Ruby interpreter offer such functionalities without any modification of the Ruby interpreter

Observation Phase

In practice, Ruby on Rails allows to write plugins for a given application. We wrote a plugin called *trace_vars* to trace and log the state of the application at each executed instruction. This plugin, written in Ruby, is independent of the application source code. It is loaded when the application starts. In an application, all controller classes inherit from an *ApplicationController* class. Thus, the behaviour of the whole application can be impacted by modifying the behaviour of the *ApplicationController* class. The goal of the *trace_vars* plugin is to interrupt the execution of the interpreter at each instruction executed in a Controller that inherits from the *ApplicationController* class, and to log the application state in a trace. For that purpose, we use the *set_trace_func* function provided by the Ruby interpreter for debugging purposes. In our implementation, we use this function to intercept the execution of each instruction in the Ruby interpreter. At each instruction step we log the variables that are in the set of tainted variables previously defined.

Observing the Program State.

In order to obtain the set of variables that are available in the context of the program at a given execution point, we can use the metaprogramming functionalities of Ruby. In particular we can obtain all the instance variables and the local variables of the current executing ruby program by using the instructions *eval("instance_variables", binding)* and *eval("local_variables", binding)*.

Tainted Variables Identification.

In order to know if a variable in the state of the application must be logged, we can build its causality cone as defined in Section Reducing the Set of Variables. This was the approach used in (Sarrouy et al., 2009). However, the Ruby interpreter offers an interesting functionality

that is available for any object at execution: the notion of *tainted* variables. For each object, a method *tainted?* permits to know if an object is causally dependent on the external inputs of the application. This functionality allows knowing if a variable is in the set of tainted variables without having to explicitly compute the dependencies between variables. However, this functionality is only partly implemented in Ruby, as an object can have attributes that are tainted without the object itself being tainted. As a consequence, when we investigate if a variable is tainted, we have to inspect if any of its attributes is tainted or not. If one of its attributes is tainted, we log this attribute.

Naming Variables in the Log.

To generate invariants on different instances of variables that evolve in time, we must modify the common use of Daikon (the tool that computes the invariants on the traces) by providing it with the values of the variables at different points in time. To offer such functionality, we could introduce additional variables by adding the execution point of the executed instruction to the name of the logged variables. With this approach, we would have the naming rule: *programCounter VariableName*. However, we must not forget that we will have to instrument the application source code to check for the invariants. As a consequence, we must find a way to link the program counter with the line in the source code where the invariant must be checked. To fulfill this constraint, we define the *programCounter* as a couple (*line*, *numberOfOccurrences*) where *line* is the line executed in the source code, and *numberOfOccurrences* the number of times that this line has been executed. The naming rule of variables becomes: *line_numberOfOccurrences_variableName*. During the observation phase, the variables are stored in a file that is formatted according to the grammar that can be recognized by the Daikon tool. In practice, for each execution of an action, we log the variables in a file. Then, when enough observations have been made, all files are merged into a single Daikon file. This file is in turn used to compute the invariants. An example of a trace obtained is given in Figure 7.

```
variable 17_1_@_params["user"]["login"]
  var-kind variable
  dec-type String
  rep-type java.lang.String
  enclosing-var 17_1_@_params["user"]
  comparability 0
variable 19_2_@_session["user"].@attributes["login"]
  var-kind variable
  dec-type String
  rep-type java.lang.String
  enclosing-var 19_2_@_session[user].@attributes
  comparability 0
...
17_1_@_params["user"]["login"]
"evette"
1
19_2_@_session["user"].@attributes["login"]
"evette"
1
```

Figure 7. Trace handled by Daikon

Time Window Handling.

In order to be able to find invariants that link variables of two consecutive requests, we need to keep in memory a window of two request executions. When we log the variables for a given action execution, we actually log the previous execution too. In practice, there can be numerous paths that can lead to a given action execution. In order to be able to discover all relations between a given action and its predecessors, we must execute all possible previous actions. This complicates the learning phase, but it is a testing problem that can be easily solved.

Generating Invariants

The invariants are generated with Daikon, a tool commonly used by the software testing community. Daikon is able to generate from a trace file likely invariants on the values held by the variables. The variable values must have been logged in the trace file. Considering the trace exhibited by Figure 7, one can see that the user login parameter of the request at line 17 of the program is equal to the user login in the session at line 19 of the program. The value held by this variable at line 19 corresponds to the second time line 19 is executed. This side effect comes from the fact that there are in practice two instructions at this line. The first one is a call to a method, and the second corresponds to an assignment of the result returned by the previous method into the considered variable. This complexity is due to the fact that we are working at source code level and not at a bytecode or binary level (with atomic instructions). The invariants that we try to discover are those proposed in Section Classes of Invariants. From the file generated for the UsersController and login action, we obtain the invariant shown by Figure 8. This invariant is exactly the one pointed out in Section **Invariant Constraints in an Application**.

```
17_1_@_params["user"]["login"] == 19_2_@_session[:user][:login]
```

Figure 8. Example of Invariant generated by Daikon

Run-Time Invariants Checking




Figure 9. Generation of invariant checking

In order to check the invariants at run-time, the application is instrumented as illustrated by Figure 9. As previously shown, the invariants have semantically all the pieces of information required to weave them in the application source code. Indeed, each invariant includes the name of the variable and also the line in the source code it is related to. In order to weave the invariants in the application source code, two types of instructions need to be inserted:

- Saving the values held by the variables of interest;

- Checking the invariant itself at the right line (e.g., for the invariant of Figure 8, the invariant shall be woven at line 19 of the concerned Controller).

```
def login
  if request.post?
    # whole user is stored in the session
    if session[:user] =
      User.authenticate(params[:user][:login],
                       params[:user][:password])

      flash[:notice] = 'successful login.'
      if session[:user].admin
        redirect_to :controller => '/admin/home',
                    :action => 'index'
      else
        redirect_to :controller => '/user/home',
                    :action => 'index'
      end
    else
      flash.now[:error] = 'Login failed'
    end
  end
end
```

Figure 10. Initial Ruby Code

If we consider an excerpt of the UsersController file that concerns only the login action Figure 10, and after having woven the invariant shown by Figure 8, we obtain the listing presented in Figure 11. In this listing the instruction saving the values held by variables of interest has been inserted and the checking of the invariant is done after the call to the authenticate method. As a result, the SQL injection described in Section **Invariant Constraints in an Application** can be detected, since it would violate this invariant. Consequently, an alert is generated (see Figure 12) describing precisely the violation observed with respect to the application state.

```
def login
  if request.post?
    if isDefined?('@_params["user"]["login"]) #AutoGenerated
      IDSAsserts.store('17_1_@_params["user"]["login"]',
                      @_params["user"]["login"])
    end
    # whole user is stored in the session
    if session[:user] = User.authenticate(
      params[:user][:login],params[:user][:password])
    if isDefined?('@_session[:user]["login"]) #AutoGenerated
      IDSAsserts.store('19_2_@_session[:user][:login]',
                      @_session[:user][:login])
      assertEqualVar('17_1_@_params["user"]["login"]',
                    '19_2_@_session[:user][:login]')
    end
    flash[:notice] = 'You have been successfully logged in.'
    if session[:user].admin
      redirect_to :controller => '/admin/home',
                  :action => 'index'
    else
      redirect_to :controller => '/user/home',
                  :action => 'index'
    end
  end
end
```



```
else
  flash.now[:error] = 'Login failed'
end
end
end
```

Figure 11. Ruby instrumented code

Alerts Format

```
<Alert>
<To>
<Request>http://localhost:3000/login</Request>
<Controller>users</Controller>
<Action>login</Action>
<Invariant>
  <Description>
    17_1_@_params["user"]["login"] ==
    19_2_@_session[:user][:login]
  </Description>
  <Line>19</Line>
  <Value1>'or '1'='1</Value1>
  <Value2>secret</Value2>
  <Datetime>Mon Mar 7 17:54:45 +0200 2011</Datetime>
</Invariant>
</To>
</Alert>
```

Figure 12. Alert Example

As shown in Figure 12, an alert is an XML structure that combines two nodes: *<To>* and *<From>*. The *<To>* node concerns the currently executed request, and the *<From>* node describes the previous request. The *<From>* node is filled with information if an invariant dealing with a relationship between two consecutive actions is violated. Each of these nodes is filled with all the information required to diagnose the problem: the http request, the controller that has been called, the action that has been performed, and a description of the violated invariant with the values of the involved variables.

Preliminary Results

Even a small application (around 1000 lines of Ruby code) and a small class of possible invariant types can lead to a large number of invariants (around 10000). These invariants are woven in the source code. As a result, the small source code files containing the controllers of the application can become huge (see Table 1 for a subset of the application controllers) and the execution time of these controllers increases. We measured execution time for a given set of controller executions. The mean execution time of a controller for non-instrumented source code was about 16.5ms. The mean time for the same executions of the instrumented application reaches 125ms. As a consequence, the instrumented code is about 8 times slower than the non-instrumented one. This is clearly an important difference. However, it must be noted that it remains acceptable for a normal user as the response of the server still seems instantaneous.

Table 1

Instrumentation results

| Subset of <i>Insecure</i> files | Initial number of lines | Lines of instrumented code | Woven Invariants |
|---------------------------------|-------------------------|----------------------------|------------------|
| Application_controller.rb | 20 | 20 | 0 |
| Home_controller.rb | 11 | 537 | 239 |
| Products_controller.rb | 26 | 1503 | 675 |
| Reviews_controller.rb | 37 | 1547 | 691 |
| Users_controller.rb | 51 | 1681 | 771 |

IDS Assessment

This section presents the experiments carried out to assess the efficiency of IDS to detect potential attacks, through the assessment of the false negative rate (percentage of attacks seen as valid inputs by the IDS), and the false positive rate (percentage of false alarms).

The evaluation of these measures requires the activation of the application and the IDS under test with a *workload* corresponding to attack-free traffic representative of the various services delivered to the users, combined with a set of malicious traffic (referred to as *attackload*) that should cover various potential attack scenarios.

Sections Workload Definition and Attackload Definition discuss how the workload and the attackload were generated, respectively. Section Experiments and Results describes the experiments and the results obtained considering the *Insecure* application as a case study. The main conclusions are discussed in Section Discussion.

Workload Definition

The workload can be generated with different methods: 1) replay of execution traces collected during “normal” browsing of the web application, 2) use of synthetic traffic generators and 3) real-time browsing of the application. For our experiments, we developed a synthetic traffic generator which offers flexibility to automatically navigate through the different URL links that can be visited from the application under test and to generate various navigation scenarios.

Several web crawlers exist but most of them cannot go through web pages that require prior authentication by providing a valid (login/password) or some specific inputs. Accordingly parts of the application may not be visited. Furthermore, additional functionalities are generally needed to facilitate the generation of attack-free traffic scenarios, such as defining (i) the maximum number of pages to be visited, or (ii) the maximum duration of a user interaction session, (iii) the number of users that can simultaneously interact with the application, or (iv) the time between consecutive user interactions with the application (usually called *think time*). As the crawlers we studied did not provide such a flexibility in their configuration, we developed our own crawler (See (Akrouf, 2012) for more details).

Attackload Definition

Attackload corresponds to malicious traffic that is aimed at exploiting vulnerabilities of the application. These malicious requests can be generated by hand, or automatically by web *vulnerability scanners* and attack tools.

Many vulnerability scanners have been developed in the last years (Bau, Bursztein, Gupta, & Mitchell, 2010), (Doupé, Cova, & Vigna, 2010)). To detect vulnerabilities, they generate specially crafted inputs and analyze the corresponding responses to determine whether the input parameters submitted through the injection point are sanitized or not.

In our experiments, we used *Wasapy*, a vulnerability scanner developed at LAAS, which allows the automated detection and exploitation of different types of web vulnerabilities including SQL injections, OS commanding, File Include, XPath. The algorithm implemented in *Wasapy* is briefly presented in the following considering the example of SQL injections (for more details, see (Dessiatnikoff, Akrouf, Alata, Kaâniche, & Nicomette, 2011)).

The vulnerability detection and exploitation algorithm implemented in *Wasapy* is based on: i) the automatic generation of inputs based on a grammar that is specific to targeted vulnerabilities, and ii) the classification of the corresponding responses using data clustering techniques.

Three sets of requests are submitted at each injection point:

- R_r corresponds to requests with randomly chosen data. They are very likely to generate error pages.

- R_{ii} corresponds to syntactically invalid SQL injections that are designed to lead to unsuccessful executions.

- R_{vi} corresponds to syntactically valid SQL injections. The main issue is to automatically determine whether they lead to a *successful execution* page or to a failed execution page. To do so, these responses are compared to those associated to sets R_r and R_{ii} using a similarity distance and a data clustering algorithm.

We note S_r , S_{ii} and S_{vi} the responses associated to R_r , R_{ii} and R_{vi} respectively. The R_{vi} requests whose responses are not similar to any of the responses from S_{ii} and S_r are considered successful SQL injections.

Experiments and Results

To evaluate the IDS, we have considered as a case study the *Insecure* application presented in Section **Context and Case Study** which emulates an e-commerce web site. The users can perform atomic operations such as login to the web site, create an account, browse the list of products, write a review, add a product to a shopping cart, pay, etc. The application contains a few vulnerabilities including SQL injection, XSS, CSRF (*Cross Site Request Forgery*) and FileUpload vulnerabilities.

We have done two sets of experiments. In the first set, only non malicious traffic was submitted to the IDS to assess the rate of false positives under normal use conditions. The workload is generated automatically using the crawler discussed in Section Workload Definition. In the second set of experiments, we used *Wasapy* to identify the vulnerabilities of the application under test and automatically generate attacks that successfully exploit the identified vulnerabilities. These attacks are complemented by specially crafted attacks, not currently

covered by *Wasapy*, that are run manually. During this second set of experiments, both false negatives (lack of detection) and false positives in the presence of attackload only, are identified.

In the following, we provide more details about the two sets of experiments, and the results obtained from these experiments. The experiments were conducted on a 3.6Ghz Pentium 4 with 6GB of RAM running Linux 2.6.32. The server was running Ruby (version 1.8) Rails (version 2.3.2) and libmysql-ruby 1.8.

Experiments with workload only.

Table 2

Number of invalidated invariants: workload only

| Number of requests | Number of raised alarms | Number of invalidated invariants |
|--------------------|-------------------------|----------------------------------|
| 1623 | 20 | 4 |

Our crawler was run during four hours, considering two different classes of users: regular users (80% of the traffic) and admin user (20% of the traffic). Each request is tagged with an identifier to facilitate the processing of the alarms raised by the IDS. Table 2 summarizes the results of these tests.

20 alarms were raised corresponding to the invalidation of four distinct invariants. These correspond to some normal browsing scenarios that were not activated during the learning phase. For example, one of these invariants asserts that the confirmation password is equal to the password, which in practice is an incorrect invariant. Another one asserts that the credit card of a user who purchases products is always the same, which is also not generally the case in practice.

Note that the number of raised alarms is less relevant than the number of invalidated invariants. Indeed, the number of raised alarms may be very different from a workload execution to another: if the crawler executes n times the scenario in which the password is different from the confirmation password, n alarms will be raised that correspond to a single invalidated invariant.

Experiments with attackload only.

Table 3

Attackload only

| Number of crafted requests | Number of successful attacks | Number of raised alarms | False positives | False negatives |
|----------------------------|------------------------------|-------------------------|-----------------|-----------------|
| 13320 | 11 | 11 | 0 | 0 |

For these experiments, attackload was first automatically generated by *Wasapy*. Additionally, some attack scenarios that are not currently supported by *Wasapy* were manually generated.

The clustering algorithm of *Wasapy* was parameterized with 30 requests for each set R_r , R_{ii} and R_{vi} . Thus, 90 crafted requests are sent through each injection point and each attack category. These values are determined empirically, however more requests could also be generated. When the four categories of vulnerabilities implemented in *Wasapy* are investigated (SQL Injection, XPATH injection, File Include and Os Commanding), 90×4 , i.e., 360 crafted requests are sent for each injection point. Table 3 summarizes the results of these tests. *Wasapy* identified 37 injection points and as a consequence, sent 37×360 (i.e., 13320) crafted requests.

The IDS detected all the requests that actually succeeded in exploiting the vulnerabilities. Indeed, among the 90×37 requests sent by *Wasapy* to detect SQL vulnerabilities, only 11 were successful. The raised alarms exactly correspond to these 11 cases. There are no false positives, neither false negatives. The other requests sent on each injection point, corresponding to XPATH, OS Commanding and File Include attacks, were not successful because such vulnerabilities do not exist in the *Insecure* application. Interestingly, the IDS did not raise any alarm in these cases too.

Three additional attack scenarios were manually executed to exploit some other vulnerabilities: 1) the manual modification of the session cookie (enabling a user to log into the application without supplying any valid login/password, by using the session cookie of an already authenticated user), 2) the manual modification of requests parameters that enables to modify the price of a product during a purchase session and 3) the upload of a Javascript File instead of an image file. The manual modification of the price of a product is correctly detected as it corresponds to the violation of an invariant identified during the learning phase. However, the other attacks were not detected. The upload of a JavaScript File cannot be detected with the proposed IDS because the target of the attack is not the web application itself but the browser (it is a form of XSS attack). As a consequence, it seems difficult to be associated to an invariant. A similar observation can be made concerning the attacks based on the modification of the session cookie. It seems quite difficult to identify one invariant that could be relevant for this kind of attack.

Discussion

The experiments described in this section show that the intrusion detection system is particularly efficient to detect successful attacks on the Web application, i.e., crafted requests that actually lead to the violation of the application state integrity. No false positives and no false negatives were generated when the attackload only was sent to the application. Let us emphasize that our vulnerability detection algorithm, sending a large number of crafted requests per injection point, is particularly adapted to evaluate the number of false positives. Regarding the false negatives, some additional experiments still have to be carried out because the *Insecure* application does not include many vulnerabilities. The false negatives must be considered carefully and should be confirmed by testing the IDS on other Ruby applications. Finally, our experiments with attack free workload also illustrate the importance of the learning phase for the intrusion detection system. These experiments raised a few false positives that are associated to four invalidated invariants. These incorrect invariants were due to the learning phase, which was incomplete. A more complete learning phase would have not produced these invariants.

Conclusion

The application level anomaly based intrusion detection mechanisms presented in this paper rely on the modelling of the application normal behaviour using invariants. The violation of an invariant is then considered as an indication of an attack against the application. A first step consists in learning the normal behaviour of the application by logging at runtime the value of variables and automatically generating invariants on these data items. The invariants are then woven in the source code to detect a violation of the application normal state during its execution. Although we are looking for a reduced set of invariant classes, the experience proved us that the approach is able to discover a large number of invariants, even on a small application. Moreover, the assessment phase proved that common attacks, such as SQL injections were successfully and accurately detected by our approach. However, some attacks such as Cross Site Scripting or Session Hijacking are not detected by our approach. This is not surprising, as these attacks either impact the client browser, or do not violate the server normal application state. These attacks are thus not the type of attacks we claim to detect.

References

- Abadi, M., Budiu, M., Erlingsson, U., & Ligatti, J. (2005). *Control-flow integrity*. Paper presented at the 12th ACM conference on Computer and communications security.
- Akritidis, P., Cadar, C., Raiciu, C., Costa, M., & Castro, M. (2008). *Preventing Memory Error Exploits with WIT*. Paper presented at the IEEE Symposium on Security and Privacy.
- Akrout, R. (2012). *Analyse de vulnérabilités et évaluation de systèmes de détection d'intrusions pour les applications Web*. (12747)
- Bau, J., Bursztein, E., Gupta, D., & Mitchell, J. (2010). *State of the art: Automated black-box web application vulnerability testing*. Paper presented at the IEEE Symposium on Security and Privacy, Oakland, USA.
- Castro, M., Costa, M., & Harris, T. (2006). *Securing software by enforcing data-flow integrity*. Paper presented at the 7th USENIX Symposium on Operating Systems Design and Implementation.
- Cavallaro, L., & Sekar, R. (2008). Anomalous Taint Detection: Secure Systems Laboratory, Stony Brook University.
- Cova, M., Balzarotti, D., Felmetzger, V., & Vigna, G. (2007, September). *Swaddler: An Approach for the Anomaly-based Detection of State Violations in Web Applications*. Paper presented at the International Symposium on Recent Advances in Intrusion Detection (RAID), Gold Coast, Australia.
- d'Ausbourg, B. (1994). *Implementing secure dependencies over a network by designing a distributed security subsystem*. Paper presented at the Third European Symposium on Research in Computer Security (ESORICS'94).
- Dessiatnikoff, A., Akrou, R., Alata, E., Kaâniche, M., & Nicomette, V. (2011). *A Clustering Approach for Web Vulnerabilities Detection*. Paper presented at the 17th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC-2011), Pasadena, CA, USA.
- Doupé, A., Cova, M., & Vigna, G. (2010). *Why Johnny can't pentest : An analysis of black-box web vulnerability scanners*. Paper presented at the DIMVA 2010.
- Ernst, M. D., Perkins, J. H., Guo, P. J., McCamant, S., Pacheco, C., Tschantz, M. S., & Xiao, C. (2007). The Daikon system for dynamic detection of likely invariants. *Science of Computer Programming*, 69, 35-45.
- Felmetzger, V., Cavedon, L., Kruegel, C., & Vigna, G. (2010, August). *Toward Automated Detection of Logic Vulnerabilities in Web Applications*. Paper presented at the 19th USENIX Security Symposium, Washington, DC.
- IBM. (2012). X-Force 2012 Mid-year Trend and Risk report, from <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=WGL03014USEN>
- Karr, M. (1976). *Affine relationships among variables of a program*. Paper presented at the Acta Informatica.

- Kiriansky, V., Bruening, D., & Amarasinghe, S. (2002). *Secure execution via program shepherding*. Paper presented at the Usenix Security Symposium.
- Ludinard, R., Totel, E., Tronel, F., Nicomette, V., Kaâniche, M., Alata, E., . . . Bachy, Y. (2012). *Detecting Attacks against Data in Web Applications*. Paper presented at the International Conference on Risks and Security of Internet and Systems (CRiSIS 2012), Cork, UK.
- Robertson, W. K., Vigna, G., Kruegel, C., & Kemmerer, R. A. (2006, February). *Using Generalization and Characterization Techniques in the Anomaly-based Detection of Web Attacks*. Paper presented at the Network and Distributed System Security Symposium (NDSS 2006), San Diego, CA.
- Sarrouy, O., Totel, E., & Jouga, B. (2009). *Application data consistency checking for anomaly based intrusion detection*. Paper presented at the 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS 2009), Lyon, France.
- Vigna, G., Robertson, W., Kher, V., & Kemmerer, R. A. (2003, December). *A Stateful Intrusion Detection System for World-Wide Web Servers*. Paper presented at the Annual Computer Security Applications Conference (ACSAC 2003), Las Vegas, NV.