JavaScript Instrumentation in Practice

Haruka Kikuchi[1], Dachuan Yu[2], Ajay Chander[2], Hiroshi Inamura[2], and Igor Serikov[2]

[1] NTT DOCOMO, Inc.
[2] DOCOMO Communications Laboratories USA, Inc.

Abstract. JavaScript has been exploited to launch various browser-based attacks. Our previous work proposed a theoretical framework applying policy-based code instrumentation to JavaScript. This paper further reports our experience carrying out the theory in practice. Specifically, we discuss how the instrumentation is performed on various JavaScript and HTML syntactic constructs, present a new policy construction method for facilitating the creation and compilation of security policies, and document various practical difficulties arose during our prototyping. Our prototype currently works with several different web browsers, including Safari Mobile running on iPhones. We report our results based on experiments using representative real-world web applications.

1 Introduction

The success of the Web can be contributed partly to the use of client-side scripting languages, such as JavaScript [3]. Programs in JavaScript are deployed in HTML documents. They are interpreted by web browsers on the client machine, helping to make web pages richer, more dynamic, and more "intelligent."

As a form of mobile code, JavaScript programs are often provided by parties not trusted by web users, and their execution on the client systems raises security concerns. The extent of the problem [2,14] ranges from benign annoyances (*e.g.,* popping up advertisements, altering browser configurations) to serious attacks (*e.g.,* XSS, phishing).

In previous work [19], we formally studied the application of program instrumentation to enforcing security policies on JavaScript programs. Specifically, we clarified the execution model of JavaScript (particularly, higher-order script—script that generates other script at runtime), presented its instrumentation as a set of syntactic rewriting rules, and applied edit automata [12] as a framework of policy management. Focusing on articulating the generic instrumentation techniques and proving their correctness, the previous work is necessarily abstract on several implementation aspects, and can be realized differently when facing different tradeoffs. For example, among other possibilities, the instrumentation process can be carried out either by a browser on the client machine or by a proxy sitting on the network gateway.

In this paper, we discuss the practical application of the above theory to real-world scenarios. Specifically, we have completed a relatively mature prototype following a proxy-based architecture. In this prototype, both the instrumentation process and the policy input are managed by a proxy program, which is situated on the network gateway (or an enterprise firewall) and is maintained separately from the JavaScript programs of

G. Ramalingam (Ed.): APLAS 2008, LNCS 5356, pp. 326–341, 2008.
© Springer-Verlag Berlin Heidelberg 2008

concern. The browser is set up to consult the proxy for all incoming traffic, and the proxy sends instrumented JavaScript programs (and HTML documents) to the browser.

Such an architecture provides centralized interposition and policy management through the proxy, thus enabling transparent policy creation and update for the end user. It requires no change to the browser implementation; therefore, it is naturally applicable to different browsers, modulo a few browser-specific implementation issues. As part of our experiments, we applied our proxy to the Safari Mobile browser on an iPhone [1]. The instrumentation and security protection worked painlessly, even though no software or browser plug-in can be installed on an iPhone. Furthermore, the proxy-based architecture poses minimal computation requirement on the client device for securely rendering web pages (as shown in our experiments, some popular web pages contain several hundred kilobytes of JavaScript code; their instrumentation takes a nontrivial amount of time). It is thus suitable for deployment with the use of mobile browsers.

We have successfully run our prototype with several browsers: Firefox, Konqueror, Opera, Safari, and Safari Mobile on an iPhone [1]. IE is currently only partially supported, because its behaviors on certain JavaScript code are significantly different than those of the other tested browsers. For the experiments, we applied a selected set of policies, which triggered the instrumentation of a variety of JavaScript constructs. We hand-picked some representative web pages as the instrumentation target. Measurements are made on both the proxy overhead and the client overhead. The initial numbers matched our expectations, noting that our prototype had not been optimized for performance.

2 Background

2.1 Previously: A Theoretical Framework

The generic theoretical framework for JavaScript instrumentation is illustrated in Figure 1 (reused from previous work [19]). Beyond the regular JavaScript interpreter provided in a browser, three modules are introduced to carry out the instrumentation and policy enforcement. A rewriting module ι serves as a proxy between the browser and the network traffic. Any incoming HTML document D, possibly with JavaScript code embedded, must go through the rewriting module first before reaching the browser.

The rewriting module identifies security-relevant actions A out of the document D and produces instrumented code $check(A)$ that monitors and confines the execution

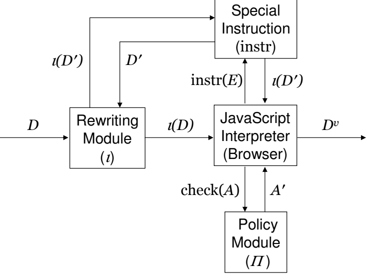

Fig. 1. JavaScript instrumentation for browser security

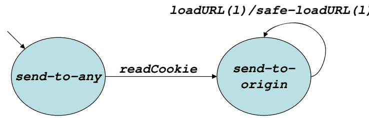

Fig. 2. Example edit automaton for a cookie policy

of A. This part is not fundamentally different from conventional instrumentation techniques [18,4]. However, this alone is not sufficient for securing JavaScript code, because of the existence of higher-order script (*e.g.,* $document.write(E)$, which evaluates E and use the result as part of the HTML document to be rendered by the browser)—some code fragments (*e.g.,* those embedded in the value of E) may not be available for inspection and instrumentation until at runtime. To address this, the rewriting module wraps higher-order script (*e.g.,* E) inside a special instruction (*e.g.,* $instr(E)$), which essentially marks it to be instrumented on demand at runtime.

With the rewriting module as a proxy, the browser receives an instrumented document $\iota(D)$ for rendering. The rendering proceeds normally using a JavaScript interpreter, until either of the two special calls inserted by the rewriting module is encountered. Upon $check(A)$, the browser consults the policy module Π, which is responsible for maintaining some internal state used for security monitoring, and for providing a replacement action A' (a secured version of A) for execution. Upon $instr(E)$, the implementation of the special instruction $instr$ will evaluate E to D' and invoke the rewriting module at runtime to perform the necessary rewriting on higher-order script. The instrumented $\iota(D')$ will then be sent back to the browser for rendering, possibly with further invocations of the policy module and special instruction as needed.

The policy module Π essentially manages an edit automaton [12] at runtime for security enforcement. A simple policy is illustrated in Figure 2, which restricts URL loading to prevent potential information leak after cookie is read. Following this, the policy module updates the automaton state based on the input actions A that it receives through the policy interface $check(A)$, and produces output actions A' based on the replacement actions suggested by the automaton. For example, if the current state is `send-to-origin`, and the input action is `loadURL(l)`, then the current state remains unchanged, and the output action becomes `safe-loadURL(l)`, which performs necessary security checks and user promptings before loading a web page.

2.2 This Paper: A Proxy-Centric Realization

Focusing on the formal aspects and correctness of policy-based instrumentation on higher-order script, the previous work is largely abstract on the implementation aspects. Subsequently, we have conducted thorough prototyping and experiments on the practical realization of the formal theory. Several interesting topics were identified during the process, which we believe will serve as useful contributions to related studies.

Overall, we have opted for a proxy-centric realization for its browser independence and low computation overhead on client devices. This can be viewed as an instantiation of the framework in Figure 1 for a specific usage scenario—the use with multiple

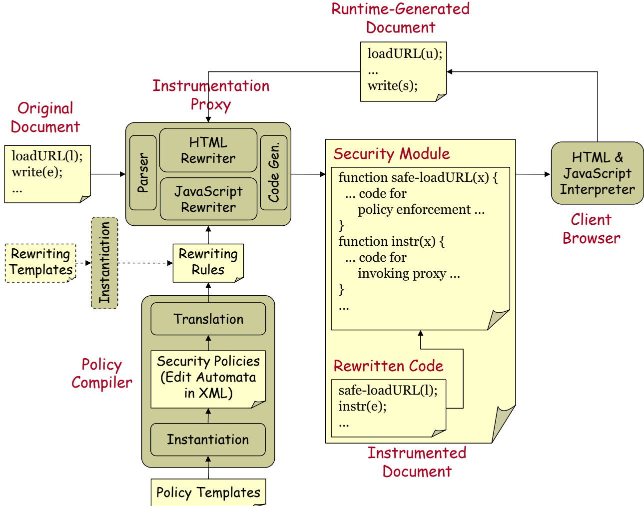

Fig. 3. A proxy-centric realization

mobile browsers. The instantiated architecture is more accurately depicted in Figure 3. The basic requirement here is that the proxy takes care of the rewriting and policy input, without changing the browser implementation. As a result, the tasks of the policy module Π and the special instruction $instr$ in Figure 1 have to be carried out in part by regular JavaScript code (the Security Module in Figure 3). Such JavaScript code is inserted into the HTML document by the proxy based on some policy input.

The rewriting process is carried out on the proxy using a parser, two rewriters, and a code generator. The rewriters work by manipulating abstract syntax trees (ASTs) produced by the parser. Transformed ASTs are converted into HTML and JavaScript by the code generator before fed to the browser that originally requested the web pages.

The rewriting rules direct the rewriters to put in different code for the security module when addressing different policies. On the one hand, the rewriters need to know what syntactic constructs to look for in the code and how to rewrite them. Therefore, they require low-level policies to work with syntactic patterns. On the other hand, human policy designers should think in terms of high-level actions and edit automata, not low-level syntactic details. We use a policy compiler to bridge this gap, which is a stand-alone program that can be implemented in any language.

We now summarize three major technical aspects of this proxy-centric realization. These will be expanded upon in the next three sections, respectively.

The first and foremost challenge is raised by the flexibility of the JavaScript language. Much effort is devoted to identifying security-relevant actions out of JavaScript code—there are many different ways of carrying out a certain action, hence many different syntactic constructs to inspect. We introduce a notion of *rewriting templates* to help managing this. Policy designers instantiate rewriting templates to rewriting rules, which in turn guides the instrumentation. This allows us to stick to the basic design without being distracted by syntactic details.

Next, the rewriting templates are still too low-level to manage. In contrast, the theoretical framework used edit automata for policy management. We support edit automata using an XML representation, and compile them into rewriting rules. As a result, we can handle all policies allowed by the theoretical framework. Nonetheless, edit automata are general-purpose, and browser-based policies could benefit further from domain-specific abstractions. We have identified some useful patterns of policies during our experiments. We organize these patterns as *policy templates*, which are essentially templates of specialized edit automata. We again use an XML representation for policy templates, which enables a natural composition of simple templates to form compound ones.

Finally, using a proxy, we avoid changing the browser implementation. This provides platform independence and reduces the computation overhead on the client devices. As a tradeoff, there are some other difficulties. In particular, the interfaces to the special instruction $instr$ and the policy module Π have to be implemented in regular JavaScript code and inserted into the HTML document during instrumentation. Such code must be able to call the proxy at runtime, interact with the user in friendly and meaningful ways, and be protected from malicious modifications. Furthermore, different browsers sometimes behave differently on the same JavaScript code, an incompatibility that we have to address carefully during instrumentation.

3 Rewriting JavaScript and HTML

We now describe the details of the instrumentation, which is carried out by transforming abstract syntax trees (ASTs). Specifically, the proxy identifies pieces of the AST to be rewritten, and replaces them with new ones that enforce the desired policies. Although a clean process in Figure 1, action replacement in the actual JavaScript language raises interesting issues due to the flexible and dynamic nature of the language. Specifically, actions A in Figure 1 exhibit themselves in JavaScript and HTML through various syntactic forms. Some common examples include property access, method calls, and event handlers. Therefore, upon different syntactic categories, the instrumentation should produce different target code, as opposed to the uniform $check(A)$. We specify what syntactic constructs to rewrite and how to rewrite them using rewriting rules.

We summarize commonly used rewriting rules for various syntactic constructs in Table 1 as *rewriting templates*. The first column shows the template names and parameters. The second shows the corresponding syntactic forms of the code pieces to be rewritten (which correspond to actions A). The third shows the target code pieces used to replace the original ones (which correspond to $check(A)$). These templates are to be instantiated using relevant JavaScript entities. Examples are given in the last column.

For example, the first row $(\texttt{Get}, obj, prop)$ specifies how to rewrite syntactic constructs of reading properties (both fields and methods). `Get` is the name of the template, and *obj* and *prop* are parameters to be instantiated with the actual object and property. Two sample instantiations are given. The first is on the field access `document.cookie`. Based on $(\texttt{Get}, \texttt{document}, \texttt{cookie})$, the JavaScript rewriter will look for all AST pieces of property access (*i.e.,* those of the shape *obj.prop* and $obj[\text{"}prop\text{"}]$), and replace them with a call to a *redirector* function `sec.GetProp`. Here `sec` is a JavaScript

Table 1. Rewriting templates

Rewriting templates	Sample Code patterns	Rewritten code (redirectors)	Sample instantiation
$(\texttt{Get}, obj, prop)$	$obj.prop$ $obj[\text{“}prop\text{”}]$	$\texttt{sec.GetProp}(obj, prop)$	$(\texttt{Get}, \texttt{document}, \texttt{cookie})$ $(\texttt{Get}, \texttt{window}, \texttt{alert})$
$(\texttt{Call}, obj, meth)$	$obj.meth(E, \dots)$ $obj[\text{“}meth\text{”}](E, \dots)$	$\texttt{sec.CallMeth}(obj, meth, E_i, \dots)$	$(\texttt{Call}, \texttt{window}, \texttt{open})$
$(\texttt{GetD}, dprop)$	$dprop$	$\texttt{sec.GetDProp}(dprop)$	$(\texttt{GetD}, \texttt{location})$
$(\texttt{CallD}, dmeth)$	$dmeth(E, \dots)$	$\texttt{sec.isEval}(dmeth)$? $\texttt{eval}(\texttt{sec.instrument}(E_i, \dots))$: $\texttt{sec.CallDMeth}(dmeth, E_i, \dots)$	$(\texttt{CallD}, \texttt{open})$
$(\texttt{Set}, obj, prop)$	$obj.prop = E$ $obj[\text{“}prop\text{”}] = E$ $obj.prop \mathrel{+}= E$ $obj[\text{“}prop\text{”}] \mathrel{+}= E$	$\texttt{sec.SetProp}(obj, prop, E_i)$ $\texttt{sec.SetPropPlus}(obj, prop, E_i)$	$(\texttt{Set}, \texttt{document}, \texttt{cookie})$ $(\texttt{Set}, \texttt{window}, \texttt{alert})$
$(\texttt{SetD}, dprop)$	$dprop = E$ $dprop \mathrel{+}= E$	$\texttt{sec.SetDProp}(dprop, E_i)$ $\texttt{sec.SetDPropPlus}(dprop, E_i)$	$(\texttt{SetD}, \texttt{location})$ $(\texttt{SetD}, \texttt{open})$
$(\texttt{Event}, tag, attr)$	$<tag\ attr=\text{“}E\text{”}>$	$<tag\ attr=\text{“}\texttt{sec.Event}(\texttt{this}, E_i)\text{”}>$	$(\texttt{Event}, \texttt{button}, \texttt{onclick})$
$(\texttt{FSrc})$	<img src=“U” onerror=“E”> <iframe src=“U” onload=“E”>	<img src=“ ” onerror=“setAttribute(onerror, “E_i”); sec.FSimg(this, U)”> <iframe src=“ ” onload=“setAttribute(onload, “E_i”); sec.FSiframe(this, U)”>	

object inserted by our proxy; it corresponds to the “Security Module” in Figure 3. Among other tasks, `sec` maintains a list of private references to relevant JavaScript entities, such as `document.cookie`. The body of the redirector above will inspect the parameters obj and $prop$ as needed to see if they represent `document.cookie`. If yes, the redirector proceeds to carry out a replacement action supplied during template instantiation.

The implementation of the replacement action is the topic of Section 4. For now, it suffices to understand the replacement action simply as JavaScript code. It can perform computation and analysis on the arguments of the redirector, provide helpful promptings to the user, and/or carry out other relevant tasks. One typical task that it carries out is to advance the monitoring state of the edit automaton used by the security policy.

The second example of the `Get` category is on accessing `window.alert`. Note that JavaScript allows a method to be accessed in the same way as a field. For example, `var f` $=$ `window.alert` assigns the method `window.alert` to a variable `f`. This is handled during rewriting using the same `Get` category as described above, and the body of `sec.GetProp` can monitor such access and implement related policies (*e.g.,* to replace the access to `window.alert` with the access to an instrumented version `sec.alert`).

The remainder of the table follows the same intuition. $(\texttt{Call}, obj, meth)$ tells the rewriter to look for syntactic categories relevant to method calls (*e.g.,* $obj.prop(E, \dots)$, $obj[\text{”}prop\text{”}](E, \dots)$) and produce a call to a redirector `sec.CallMeth`. The argument E to the method invocation is rewritten to E_i following the same set of rewriting rules (the same also applies to other cases in the table). $(\texttt{GetD}, \texttt{dprop})$ and $(\texttt{CallD}, \texttt{dmeth})$ are for accessing default properties and calling default methods. $(\texttt{Set}, \texttt{obj}, \texttt{prop})$ and $(\texttt{SetD}, \texttt{dprop})$ are for setting object properties and default properties.

Sometimes relevant actions are coded as event handlers of certain HTML tags. For example, the code <`button onclick` $=$ “`alert()`”> raises an alert window whenever the button is clicked. The template $(\texttt{Event}, tag, attr)$ captures such event handling constructs. The redirector `sec.Event` takes the instrumented expression E_i and the parent object `this` of the attribute as arguments. The exact implementation of the redirector

depends on the security policy. Typically, the internal state of the policy edit automaton is updated while entering and exiting the corresponding event.

The last template (`FSrc`) is on the loading of external resources, such as those initiated by <img src=“E”> and <iframe src=“E”>. Instead of directly loading such a resource, we use an event handler with a redirector for interposition. The `src` attribute is modified from the original URL U to a space character to trigger immediately a corresponding event (*e.g.,* `onerror` for `img`, `onload` for `iframe`). In the rewritten event handler, after binding the original event handler (rewritten from E to E_i) back to the event for execution later, we call a redirector function (*e.g.,* `sec.FSimg`, `sec.FSiframe`). The redirector implementation depends on the policy. For example, it may check the target domain of the URL U to identify where the HTTP request is sent, perform other URL filtering, and/or present user prompting. In general, this rule is also applicable to some other cases of HTTP requests (*e.g.,* <a href = “U”>), except different event handlers (*e.g.,* `onclick`) and redirectors (*e.g.,* `sec.FShref`) are used accordingly.

Although designed mainly for rewriting actions (A), the rewriting templates are also applicable for handling higher-order script. For example, `document.write`(E) should be rewritten using `sec.instrument` (the realization of the `instr` of Figure 1). This can be represented using the `Call` template as (`Call`, `document`, `write`). In our prototype, however, the handling of higher-order script is directly coded in the rewriter for efficiency, since it is always performed regardless of what the security policy is.

There are some other cases where built-in rewriting rules are applied. Selected ones are given in the companion technical report [9]. Most of the cases are designed to handle script that is not available statically, but rather generated or loaded at runtime. Others are for facilitating error handling and instrumentation transparency.

In summary, the proxy performs rewriting by syntax-directed pattern matching on ASTs, and redirectors are used to implement appropriate interposition logic and security polices. Built-in rewriting rules are always applied, but a policy designer may customize action replacement using rewriting templates. If the same rewriting template is instantiated multiple times on different entities, a single rewriting rule with a merged function body for the redirector is produced. The proxy only rewrites cases described by the given rewriting rules. If a certain security policy only uses one rewriting template (*e.g.,* `get`, possibly with multiple instantiations), then only one rewriting rule is produced, and only the corresponding syntactic constructs (*e.g.,* `obj.prop`, `obj[“prop”]`) are rewritten. This avoids unnecessary rewriting and execution overhead.

4 Policy Writing and Management

The rewriting templates and their redirector code in Section 3 serve as a low-level policy mechanism. It allows policy designers to focus on the abstract notion of “actions” without being distracted by the idiosyncrasies of JavaScript syntax. However, it provides little help on encoding edit automata. If used for construction, a policy designer needs to implement states and transitions of policy automata in the redirectors.

We now discuss a more manageable framework that directly accommodates the notion of edit automata, allowing policy designers to focus on implementing replacement actions (*e.g.,* insertion of runtime checks). The key of this framework is the policy compiler in Figure 3, a stand-alone program compiling policies into rewriting rules off-line.

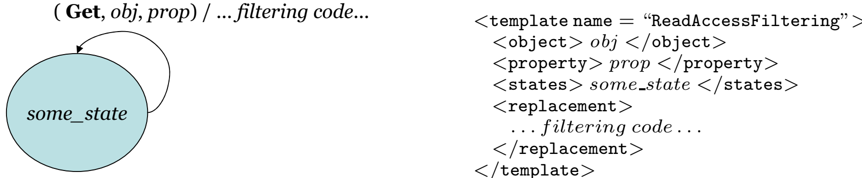

Fig. 4. Read access filtering

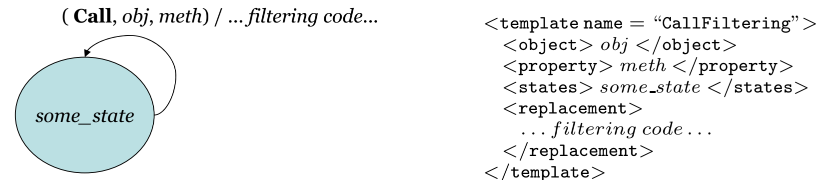

Fig. 5. Call filtering

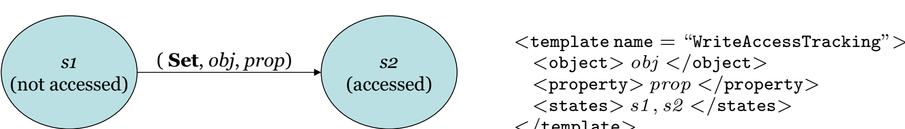

Fig. 6. Write access tracking

In essence, an edit automaton concerns a set of states (with one as an initial state), a set of actions, and a set of state transitions. In any state, an input action determines which transition to take and which output action to produce. We use an XML file to describe all these aspects for an edit automaton as a security policy. Due to space constraints, we refer interested readers to the companion technical report [9] for the details of the XML representation and its compilation to rewriting rules.

During our policy experiments, we identified several commonly used patterns of edit automata. We organize these as *policy templates*. Instead of always describing an edit automaton from scratch, a policy designer may instantiate a relevant template to quickly obtain a useful policy. We illustrate this with examples.

One pattern is on filtering read access to object properties (both fields and methods). The (fragment of) edit automaton is shown in Figure 4, together with an XML representation. Note that this is essentially a template to be instantiated using a specific object, property, state, and filtering code. It captures a common pattern where the property access is filtered without applying any state transition.

One may also filter method calls in a similar pattern (Figure 5). This is typically used to insert security checks for method calls. Another pattern is on tracking write access (Figure 6). This tracks the execution of property writing by transitioning the state of the automaton. The replacement action is the same as the input action. The corresponding policy template specifies two states, but no replacement action.

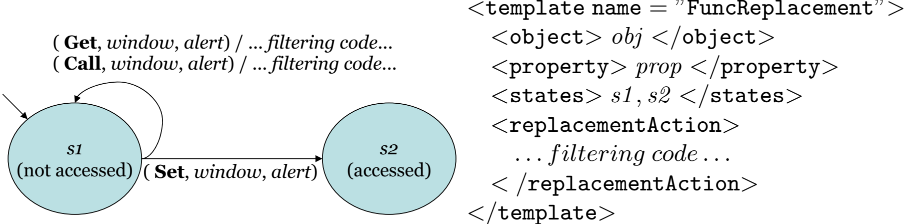

Fig. 7. Function replacement

We have implemented a total of 14 policy templates, including four on tracking property access (object and default properties; read and write access), four on filtering property access (object and default properties; read and write access), four on method calls (tracking and filtering; object and default methods), one on tracking inlined event handlers (e.g., `<body onunload="window.open();">`), and one on filtering implicit HTTP requests (e.g., images and inlined frames). The policy compiler expands instantiated policy templates into edit automata, and further compilation follows.

The above templates each describe a particular aspect of security issue. It is often the case that multiple templates are used together to enforce a useful policy. As an example, consider a simple policy of adding a prefix "`Security Module :` " to the alert text produced by `window.alert`. Naturally, we implement a replacement action `myAlert` as follows: `function myAlert(s) { window.alert("Security Module : " + s); }`.

This obviously requires some filtering on read access and calls to `window.alert`, as illustrated in Figures 4 and 5. For example, the former is applicable to rewriting code from `f = window.alert` to `f = myAlert`, and the latter is applicable to rewriting code from `window.alert("Hello")` to `myAlert("Hello")`.

An additional complication is that JavaScript code in the incoming document may choose to rewrite `window.alert` for other functionalities:

```
window.alert = function(s) {};
window.alert("a debugging message to be ignored");
```

Here, `window.alert` is redefined as an "ignore" function. It would be undesirable to perform the same filtering after the redefinition. Therefore, a more practical policy is to filter read access and calls to `window.alert` only if it has not been redefined.[1] This can be addressed using write access tracking as in Figure 6.

A compound edit automaton can be obtained by combining read access filtering, call filtering, and write access tracking, as shown on the left side of Figure 7. In our XML representation, this can be achieved using a direct combination of the three instantiated templates. This is a very common pattern, because any methods could be redefined in incoming JavaScript code. Therefore, we also support a compound template `FuncReplacement` (shown on the right side of Figure 7) to directly represent this pattern. The actual policy on `window.alert` can then be obtained by instantiating this compound template with the corresponding parameters.

[1] Readers might wonder if this could be exploited to circumvent the instrumentation. The answer is no, because the function body of the new definition would be instrumented as usual.

We have also implemented some other compound templates to represent common ways of template composition. Examples include one on filtering default function calls up to redefinition and one on tracking the invocation of global event handlers.

As a summary, instantiated policy templates are expanded into edit automata, which in turn are compiled into rewriting rules and redirectors. The XML-based representation supports arbitrary edit automata for expressiveness, and some domain-specific templates are introduced to ease the task of policy construction.

5 Runtime Interaction with Proxy and with Client

5.1 Structure of Instrumented Documents

We have described two stand-alone components—one is the proxy for the instrumentation of incoming contents at runtime, the other is a policy compiler compiling high-level security policies (edit automata) into low-level rewriting rules off-line. These two components collaborate to complete the instrumentation tasks.

The structure of an instrumented document is shown in Figure 8. Based on the rewriting rules, the proxy replaces relevant syntactic constructs in the incoming content with calls to redirectors. The proxy also inserts a *security module*, which contains the realization of the special instruction $instr$ and policy module Π in Figure 1. Specifically, $instr$ exhibits itself as a utility function `sec.instrument` (more in Section 5.2), and the policy module Π is interfaced through the redirectors. The redirectors call certain transition functions for maintaining edit automata at runtime, and invoke some replacement actions provided by policy designers. These functions may also refer to other utility functions as needed. We will discuss one such utility function on user notification in Section 5.3. Other commonly used utility functions include those for IP normalization, URL filtering, and policy object embedding and manipulation.

Since the security module is realized as regular JavaScript code, it resides in the same execution environment as incoming contents, and thus must be protected from malicious exploits. For example, malicious JavaScript code may modify the implementation of the security module by overwriting the functions of security checks. We organize the security module in a designated object named `sec`, and rename all other objects that could cause conflicts (*i.e.,* changing `sec` into `_sec`, `_sec` into `__sec`, and so on).

Although `sec` is usually inserted by the proxy during rewriting, sometimes a window could be created without going through the proxy at all. For example, the incoming

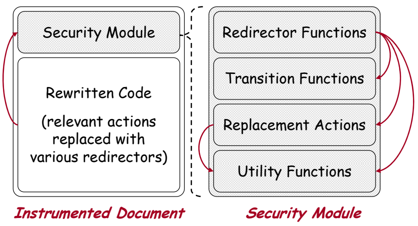

Fig. 8. Structure of an instrumented document

content may use window.open() to open up an empty new window without triggering the proxy. Once the window is open, its content could be updated by the JavaScript code in the parent window. In response, we explicitly load a sec object into the empty new window when needed, so as to correctly enforce its future behaviors.

Another potential concern is that the incoming code could modify the call sites to the redirectors through JavaScript's flexible reflection mechanisms. Suppose the rewritten document contains a node that calls sec.GetProp. Using features such as innerHTML, incoming script could attempt to overwrite the node with some different code, such as exploit. This is in fact one form of higher-order script, and it is correctly handled, provably [19], following the theoretical framework of Figure 1. In particular, the runtime-generated script exploit will be sent back to the proxy for further instrumentation before executed; therefore, it cannot circumvent the security policy. In essence, the effect of using the runtime-generated code exploit as above is not different from statically using exploit in the original content directly.

Finally, the proxy-centric architecture is browser-independent in theory, and the proxy should work for all browsers. In practice, however, different browsers sometimes behave differently when rendering the same content, due to either implementation flaws or ambiguity and incompleteness of the language specification [3] (*e.g.,* some behaviors are undefined). The companion technical report provides more details on this aspect.

5.2 Calling Proxy at Runtime

The utility function instrument needs to call the proxy at runtime to handle higher-order script. Therefore, we need to invoke the rewriters on the proxy from within the JavaScript code on the client. This is done with help of the XMLHttpRequest object [17] (or ActiveX in IE), which allows us to send runtime-generated JavaScript/HTML code to the proxy and receive their rewritten result.

An interesting subtlety is that XMLHttpRequest is restricted to communicate only with servers that reside in the same domain as the origin of the current document. We use a specially encoded HTTP request, targeting the host of the current document, as the argument to XMLHttpRequest. Since all HTTP requests go through the proxy, the proxy is able to intercept relevant requests based on the path encoding and respond with the instrumentation result. Some key code, simplified (*e.g.,* the handling of the scope chain is omitted) for ease of reading, is given below for illustration:

```
// This function sends str to the proxy for instrumentation.
// str is either HTML or JavaScript code, depending on type.
function instrument (type, str) {
  var xhr = new XMLHttpRequest();
  var url = "http://" + location.hostname + "/?__proxy__/"
                      + type + "&url = " + escape(location.href);
  try{
    xhr.open("POST", url, false);   // false specifies synchronous communication.
    xhr.send(str);
  }catch(e){...}
  return xhr.responseXML;           // The result of the instrumentation is in XML.
}
```

5.3 User Interaction

Effective user interaction upon a policy violation is important for practical deployment. A simple notification mechanism such as a dialogue box may not appear sufficiently friendly or informative to some users. Upon most policy violations, we overlay the notification messages on top of the rendered content. This better attracts the user's attention, disables the user's access to the problematic content, and allows the user to better assess the situation by comparing the notification message with the rendered content.

This would be straightforward if we were to change the browser implementation. However, in the proxy architecture, we need to implement such notification in HTML. To enable the overlaying effect, we use a combination of JavaScript and Cascading Style Sheets (CSS) to provide the desired font, color, visibility, opacity and rendering areas. An interesting issue occurs, however, because such functionality works by directly manipulating the document tree. This manipulation happens after the entire document is loaded, *e.g.,* by using an `onload` event handler. Unfortunately, a policy violation may occur before the `onload` event. In this case, we fall back to use either dialogue boxes, silent suppression, or page redirection based on different code and error scenarios [9].

6 Experiments

We have run our proxy with several browsers: Firefox, Konqueror, Opera, Safari, Safari Mobile on iPhone, and (partially) IE. During experiments, we manually confirmed that appropriate error notifications were given upon policy violations. In this section, we report some performance measurements on well-behaved web pages to demonstrate the overhead. These measurements were made mainly using Firefox as the rendering browser, with help of the Firebug [7] add-on. Specifically, we profiled JavaScript execution in target web pages without counting in certain network-relevant activities, such as those due to the loading of inlined frames and external JavaScript sources, and the communication through `XMLHTTPRequest`. Two machines were used in the experiments—one as the proxy, the other as the client. Both machines have the same configuration: Intel Pentium 4, clock rate 3.2GHz, 1.5 GB of RAM, running FreeBSD 6.2. Microbenchmarks are given in the companion technical report due to space constraints.

6.1 Macro-benchmarks

To learn how the instrumentation works under typical browsing behaviors, we run macro-benchmarks using a selected set of policies and some popular web applications.

Policy Set.3 We crafted a set of policies to serve together as the policy input to the proxy, as listed in Table 2. The policies are selected based on both relevance to security and coverage of rewriting rules. The Cookie policy warns against the loading of dynamic foreign links after a cookie access [10], helping preventing XSS. The IFrame policy warns against foreign links serving as iframe sources, helping preventing a form of phishing. The IP-URL policy disallows dynamic IP URLs so as to prevent incoming script from analysing the presence of hosts on the local network. The Pop-up policy sets a limit on the number of pop-up windows, and restricts the behaviors of unwieldy

Table 2. Policy set and coverage of rewriting cases

	`Get`	`Call`	`GetD`	`CallD`	`Set`	`SetD`	`Event`	`FSrc`
Cookie	X							
IFrame					X			X
IP-URL					X	X		X
Pop-up	X	X	X	X	X	X		
URL-Event					X	X	X	X

Table 3. Target applications and various performance measurements

	DoCoMo *(corporate)*	LinkedIn *(social)*	WaMu *(bank)*	MSN *(portal)*	YouTube *(video)*	MSNBC *(news)*	GMap *(map)*	GMail *(email)*
size before (B)	24,433	97,728	156,834	170,927	187,324	404,311	659,512	899,840
size after (B)	28,047	144,646	141,024	252,004	232,606	495,568	959,097	1,483,577
ratio	1.15	1.48	0.90	1.47	1.24	1.23	1.45	1.65
proxy time (ms)	614	1,724	2,453	3,933	4,423	7,290	10,570	14,570
time before (ms)	94	82	553	402	104	2,832	1,369	4,542
time after (ms)	143	143	688	695	167	3,719	1,783	7,390
time ratio	1.52	1.74	1.24	1.73	1.61	1.31	1.30	1.63
time diff (ms)	49	61	135	293	63	887	414	2,848

(*e.g.,* , very small/large, out-of-boundary, and respawning) pop-ups. Finally, the URL-Event policy inspects certain event handlers (*e.g.,* `onclick`) to prevent malicious code from updating target URLs unexpectedly (*e.g.,* redirection to a phishing site after a linked is clicked). These together cover all the rewriting cases of Table 1.

Target Applications and Overheads. We hand-picked a variety of web pages as the target applications. These applications and their "code" sizes (for contents that require rewriting, including JavaScript and various HTML constructs, but excluding images, etc) before and after the instrumentation are listed in the top portion of Table 3. A variety of sizes are included, ranging from about 24KB to nearly 900KB. Recall that the proxy produces rewritten code and inserts a security module. The sizes in Table 3 are about the rewritten code only. The security module is always the same once we fix the policy set. For our policy set, the security module is 16,202 bytes; it is automatically cached by the browser, thus reducing network cost.

The ratio row shows how much the code size grew after instrumentation. In most cases, the growth was less than 50%. The worst case was the inbox of GMail, where the nearly 900KB of code grew by 65%. Interestingly, the WaMu code reduced by about 10% after instrumentation. By inspection, we found out that there was a significant amount of code commented out in the WaMu page, which was removed by the proxy.

The middle row shows proxy rewriting time, calculated based on the average of 50 runs. The figures are roughly proportional to the code sizes. The proxy spent a few seconds for smaller applications but over ten seconds for bigger ones. Since there was usually multiple code chunks processed, the client side perceptual delay was alleviated, because part of the content had started rendering even before the rewriting was done.

The bottom portion is Firefox interpretation time of the code before and after rewriting. The ratio and diff columns show the proportional and absolute time increases for rendering the instrumented code. For most applications, the absolute increase is negligible with respect to user perception. For GMail, the increase is nearly three seconds.

6.2 Safari Mobile on iPhone

We briefly report our experience of web browsing using Safari Mobile on iPhone. Although there has been some fluctuation of traffic in the WLAN, the numbers we collected still seems useful in describing the perceptual overhead introduced by the proxy.

When directly loading the applications of Table 3 without using the proxy, almost all pages started showing within 7-12 seconds, and finished loading (as indicated by a blue progress bar) within 11-24 seconds. The only exception was the inbox of GMail, which took well over a minute to be rendered. When loading the applications through the proxy, almost all pages started showing within 9-14 seconds (with exceptions described below). The finish time of the loading varied from 15 seconds to over a minute.

A few special cases are worth noting. MSNBC used a special page for rendering on iPhones; it took on average about 11 seconds for the entire page to be loaded without the proxy, and 15 seconds with the proxy. For GMap, we further experimented with the address searching functionality. It took on average 17 seconds to render the resulting map without the proxy, and 34 seconds with the proxy. GMail inbox could not be rendered through the proxy, because the proxy does not currently support the CONNECT protocol [13] for tunnelling the SSL'ed login information. In contrast, our tested desktop browsers were set up to use direct connections for SSL'ed contents.

7 Related Work

There has been work applying code instrumentation to the security of machine code and Java bytecode. SFI [18] prevents access to memory locations outside of predefined regions. SASI [4] generalizes SFI to enforce security policies specified as security automata. Program Shepherding [11] restricts execution privileges on the basis of code origin, monitors control flow transfers to prevent the execution of data or modified code, and ensures that libraries are entered only through exported entry points. Naccio [6] enforces policies that place arbitrary constraints on resource manipulations as well as policies that alter how a program manipulates resources. PoET [5] applies inlined reference monitors to enforces EM policies [16] on Java bytecode. As pointed out by several studies [4,15,19], the above work is not directly applicable to JavaScript instrumentation for user-level security policies. Some notable problems include the pervasive use of reflection and higher-order script in JavaScript, the lack of flexible and usable policy management, and the difficulty of interpositioning a prototype-based object model.

A closely related work is BrowserShield [15], which applies runtime instrumentation to rewrite HTML and JavaScript. Designed mainly as a vulnerability-driven filtering mechanism to defend browser vulnerabilities prior to patch deployment, the policies of BrowserShield are mainly about vulnerability signatures, and are written directly as JavaScript functions. After initial rewriting at an enterprise firewall, rewriting logic is

injected into the target web page and executed at browser rendering time. A browser plug-in is used to enable the parsing and rewriting on the client.

In contrast, we target general user-level policies, and much of our work has been on practical policy management. Specifically, domain-specific abstractions are used for policy construction, and a policy compiler is engaged to translate such constructed policies to syntactic rewriting rules. Different syntactic categories are rewritten based on different policies. For simple policies, only one or two kinds of syntactic constructs are rewritten, although a composite policy typically requires the rewriting of more. All the rewriting (both load-time and run-time) happens on a proxy. The rewritten page interacts with the proxy at runtime using `XMLHttpRequest`. No software or plug-in is required on the client. Therefore, our architecture is naturally applicable to work with multiple web browsers, even if software/plug-in installation is not allowed.

JavaScript instrumentation has also been applied to the monitoring of client-side behaviors of web applications. Specifically, AjaxScope [8] applies on-the-fly instrumentation of JavaScript code as it is sent to users' browsers, and provides facilities for reducing the client-side overhead and giving fine-grained visibility into the code-level behaviors of the web applications. Targeting the development and improvement of non-malicious code (*e.g.,* during debugging), AjaxScope does not instrument code that is generated at runtime (*i.e.,* higher-order script). Therefore, the rewriting is simpler and, as is, not suitable as a security protection mechanism against malicious code.

8 Conclusion and Future Work

We have presented a JavaScript instrumentation prototype for browser security. We reported experiences on instrumenting various JavaScript constructs, composing user-level policies, and using a proxy to enable runtime rewriting. Our proxy enables flexible deployment scenarios. It naturally works with multiple browsers, and does not require software installation on the client. It provides a centralized control point for policy management, and poses relatively small computation requirement on the client. Although not optimized, our prototype yields promising results on the feasibility of the approach.

In the future, we plan to conduct more experiments based on real-world web browsing patterns (*e.g.,* top URLs from web searches) and improve the support on popular browsers (most notably IE). We also plan to study the instrumentation of plug-in contents. For example, VBScript and Flash are both based on the ECMAScript standard [3], thus our instrumentation techniques are also applicable.

Our proxy-based architecture does not directly work with encrypted contents (*e.g.,* SSL). Some of our tested web pages (*e.g.,* GMail) uses SSL, but only for transmitting certain data, such as the login information. If the entire page was transmitted through SSL (as is the case of many banking sites), then the proxy cannot perform rewriting. This can be addressed by either decrypting on the proxy (if a trusted path between the proxy and the browser can be established), having the browser sending decrypted content back to the proxy, or directly implementing the instrumentation inside the browser.

Although the proxy-based architecture enables flexible deployment scenarios, a browser-based implementation may also be desirable. Some of our difficulties supporting multiple browsers have been due to their inconsistent treatment on undefined

JavaScript behaviors. If implemented inside the browser, the same parsing process would be applied to both the rendering and the instrumentation, thus avoiding extra parsing overhead and problems caused by specific semantic interpretations.

References

1. Apple Inc. Safari mobile on iphone, http://www.apple.com/iphone/internet/
2. Christey, S., Martin, R.A.: Vulnerability type distributions in CVE (2007), http://cve.mitre.org/
3. ECMA International. ECMAScript language specification. Standard ECMA-262, 3rd Edition (December 1999)
4. Erlingsson, U., Schneider, F.B.: SASI enforcement of security policies: A retrospective. In: Proc. 1999 New Security Paradigms Workshop, Caledon Hills, Ontario, Canada, pp. 87–95 (September 1999)
5. Erlingsson, U., Schneider, F.B.: IRM enforcement of Java stack inspection. In: Proc. IEEE S&P (2000)
6. Evans, D., Twyman, A.: Flexible policy-directed code safety. In: Proc. 20th IEEE S&P, pp. 32–47 (1999)
7. Hewitt, J.: Firebug—web development evolved, http://www.getfirebug.com/
8. Kiciman, E., Livshits, B.: AjaxScope: a platform for remotely monitoring the client-side behavior of web 2.0 applications. In: Proc. SOSP 2007, pp. 17–30 (2007)
9. Kikuchi, H., Yu, D., Chander, A., Inamura, H., Serikov, I.: Javascript instrumentation in practice. Technical Report DCL-TR-2008-0053, DoCoMo USA Labs (June 2008), http://www.docomolabsresearchers-usa.com/~dyu/jiip-tr.pdf
10. Kirda, E., Kruegel, C., Vigna, G., Jovanovic, N.: Noxes: a client-side solution for mitigating cross-site scripting attacks. In: Proc. 2006 ACM Symposium on Applied Computing, pp. 330–337 (2006)
11. Kiriansky, V., Bruening, D., Amarasinghe, S.P.: Secure execution via program shepherding. In: Proc. 11th USENIX Security Symposium, pp. 191–206 (2002)
12. Ligatti, J., Bauer, L., Walker, D.: Edit automata: Enforcement mechanisms for run-time security policies. International Journal of Information Security 4(2), 2–16 (2005)
13. Luotonen, A.: Tunneling TCP based protocols through web proxy servers. IETF RFC 2616 (1998)
14. OWASP Foundation. The ten most critical web application security vulnerabilities (2007), http://www.owasp.org/
15. Reis, C., Dunagan, J., Wang, H.J., Dubrovsky, O., Esmeir, S.: BrowserShield: Vulnerability-driven filtering of dynamic HTML. In: Proc. OSDI 2006, Seattle, WA (2006)
16. Schneider, F.B.: Enforceable security policies. Trans. on Information & System Security 3(1), 30–50 (2000)
17. van Kesteren, A., Jackson, D.: The XMLHttpRequest object. W3C working draft (2006), http://www.w3.org/TR/XMLHttpRequest/
18. Wahbe, R., Lucco, S., Anderson, T.E., Graham, S.L.: Efficient software-based fault isolation. In: Proc. SOSP 1993, Asheville, NC, pp. 203–216 (1993)
19. Yu, D., Chander, A., Islam, N., Serikov, I.: JavaScript instrumentation for browser security. In: Proc. POPL 2007, Nice, France, pp. 237–249 (January 2007)