Reactive non-interference for the browser: extended version

Nataliia Bielova *Dominique Devriese*
Fabio Massacci *Frank Piessens*

Report CW 602, Feb 2011

Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

Reactive non-interference for the browser: extended version

Nataliia Bielova *Dominique Devriese*
Fabio Massacci *Frank Piessens*

Report CW 602, Feb 2011

Department of Computer Science, K.U.Leuven

Abstract

Given a partially ordered set (poset) of security levels, and a labelling of inputs and outputs with such levels, non-interference (or secure information flow) is a security property expressing that outputs of level l only depend on inputs that are labelled with a level smaller than l. In other words, there is no information flow from high (confidential) levels, to low (public) levels.

For web browsers, as programs that interact intensely with a variety of principals, non-interference is an interesting security property, and several authors have studied how enforcement mechanisms for it can be incorporated in a browser, usually focusing on specific scenarios such as securing the flow of information towards advertisements, or securing mashups.

In this paper, we investigate the suitability of non-interference as a replacement for the baseline security policy of a browser, the *same-origin-policy*. We propose an enforcement mechanism that can enforce non-interference with respect to a broad class of security level posets for the full browser. We prove the security and precision of the enforcement mechanism, and implement it for the Featherweight Firefox browser model. Next, we investigate what security level posets are useful in a web context, and how inputs and outputs to the browser should be labelled. Somewhat surprisingly, our analysis shows that useful policies (which approximate but improve the current same-origin-policy) can be defined without any support for declassification.

Keywords : non-interference, browser security.

Reactive non-interference for the browser: extended version

Nataliia Bielova* Dominique Devriese† Fabio Massacci*
Frank Piessens†

Abstract

Given a partially ordered set (poset) of security levels, and a labelling of inputs and outputs with such levels, non-interference (or secure information flow) is a security property expressing that outputs of level l only depend on inputs that are labelled with a level smaller than l. In other words, there is no information flow from high (confidential) levels, to low (public) levels.

For web browsers, as programs that interact intensely with a variety of principals, non-interference is an interesting security property, and several authors have studied how enforcement mechanisms for it can be incorporated in a browser, usually focusing on specific scenarios such as securing the flow of information towards advertisements, or securing mashups.

In this paper, we investigate the suitability of non-interference as a replacement for the baseline security policy of a browser, the *same-origin-policy*. We propose an enforcement mechanism that can enforce non-interference with respect to a broad class of security level posets for the full browser. We prove the security and precision of the enforcement mechanism, and implement it for the Featherweight Firefox browser model. Next, we investigate what security level posets are useful in a web context, and how inputs and outputs to the browser should be labelled. Somewhat surprisingly, our analysis shows that useful policies (which approximate but improve the current same-origin-policy) can be defined without any support for declassification.

1 Introduction

The explosive growth of Web applications such as web-based e-mail, social networking, web banking, and others has turned the Web into one of the most important software delivery platforms. The Web browser has become a virtual machine that receives and executes a variety of interactive applications from different stakeholders. Hence, one of the key security responsibilities of a browser

*University of Trento, Italy, {bielova, massacci}@disi.unitn.it
†DistriNet Research Group, KULeuven, Belgium, {dominique.devriese, frank.piessens}@cs.kuleuven.be

is to provide proper protection mechanisms to ensure that these different applications can not interfere with each other in non-authorized ways. In today's browsers, this is achieved by enforcing the *same-origin-policy*. An *origin* is a (protocol, domain name, port) triple, and restrictions are imposed on the way in which code and data from different origins can interact.

Unfortunately, this same-origin-policy is fraught with problems. Not only is it implemented inconsistently in current browsers [18], it is also ambiguous and imprecise [2], and it fails to provide adequate protection for resources belonging to the user rather than to some origin [18]. This has led to a significant amount of research proposing improvements for web browser security, ranging from specific countermeasures for holes in the same-origin-policy to proposals for new browser architectures that basically turn a browser into a service operating system. We give a brief overview of this research area in the related work section.

Of particular importance for this paper are the various proposals that have been made to base the policy enforced by a browser on *non-interference*, or *information-flow security*. A program is non-interferent if secret inputs to the program do not influence public outputs. In other words, secret inputs should not flow (directly or indirectly) to public outputs. Non-interference can be defined with respect to a more general *information flow policy*. Such a policy is a partially ordered set (poset) of security levels l. The levels can be thought of as confidentiality levels: levels higher in the poset will label more confidential information. All input channels and output channels of the program are labelled with such a security level, and the program is non-interferent if information only flows from inputs labelled l_i to outputs labelled l_o for $l_i \leq l_o$. In other words: information only flows upward, toward more confidential levels. [1]

Non-interference has been studied intensely for several decades, and a wide variety of enforcement mechanisms have been proposed. Sabelfeld and Myers [15] provide an extensive survey of static enforcement methods, and Le Guernic [8] surveys dynamic methods. Several authors have already investigated the use of secure information flow techniques in the context of a browser, for instance to secure mashup composition [11,13], or to prevent private information to flow to advertisement providers [5,12].

Very recently, Bohannon et al. [3] proposed non-interference as a candidate replacement for the same-origin-policy. They define the notion of *reactive non-interference*, an adaptation of the classic notion of non-interference to *reactive systems*, systems that perform asynchronous I/O such as web browsers. In addition, they provided a bisimulation-based proof technique to prove the soundness of enforcement mechanisms for reactive non-interference. In a later paper, Bohannon et al. [2] develop Featherweight Firefox, an extensive formalization of a web browser as a reactive system. They also provide an implementation of Featherweight Firefox.

[1] Throughout the paper, we take "lower than" to mean "lower than or equal".

1.1 Contributions of this paper

A first contribution of this paper is the development of an enforcement technique for reactive non-interference based on *secure multi-execution* [7]. We prove two major results:

- *Security* Featherweight Firefox (in fact any reactive system in the sense of Bohannon et al. [3]) is reactive non-interferent when executed under this secure multi-execution regime.
- *Precision* For inputs for which Featherweight Firefox is "well-behaved" with respect to the policy, execution under the secure multi-execution regime will not result in changes in observable behavior for an observer at any security level.

Further, we show the value of our technique for web browsers, by implementing it for Featherweight Firefox. To the best of our knowledge, our proposal is the first to enforce a general non-interference policy for the browser as a whole.

Non-interference is parameterized by an information flow policy, so an interesting question is what policies are useful in a browser: what should be the levels, and how should we assign them to inputs and outputs? Interestingly, even without support for declassification, we show that many interesting and useful policies can be enforced and we study their behavior in example scenario's. Some of the policies feature infinite security level posets, and we demonstrate how our implementation can support them. Other examples include useful policies that cannot be enforced by access control based systems, and policies that approximate the current same-origin-policy to maintain some form of compatibility with the current web. We argue that fine-grained policies are required to achieve more compatibility and discuss an extension of our technique to policies at a finer level than Featherweight Firefox input events.

In summary, the contributions of this paper are:

- The development of a provably secure and precise enforcement mechanism for reactive non-interference. Our precision results are stronger and more general than those in related work.
- The analysis of a variety of policies that can be enforced by this mechanism, thus providing evidence of the suitability of non-interference as a replacement for the current same-origin-policy.
- The implementation of this mechanism for the Featherweight Firefox browser model.

1.2 Outline of the paper

In the next section we illustrate the problem addressed in this paper.

Section 3 gives an informal overview of our approach and Section 4 provides a formal model where we prove our main precision and security results. In

Section 5, we discuss a variety of useful policies that can be enforced by our mechanism. We provide more details about our implementation for Featherweight Firefox in Section 6. Finally, we discuss related work, and conclude.

2 Problem statement

A browser interacts with a variety of web sites, and possibly executes JavaScript code downloaded from these sites. Hence, a browser should enforce some security policy to make sure that these sites do not interfere with each other in undesirable ways. Today's browsers enforce the same-origin-policy, an access-control policy where browser resources are tagged with their origin, and access to resources is limited to code coming from the same origin.

The same-origin-policy has many problems, and has been criticized by many authors [10, 18]. Some of the issues, such as for instance the fact that different browser resources use different definitions of the notion of origin, can be considered implementation bugs or inconsistencies, and they could in principle be addressed without fundamentally changing the same-origin access control policy (even though, as Singh et al. point out [18], the incompatibility burden of such fixes can be substantial). While such issues are important, they are not what this paper is about.

Other limitations of the same-origin-policy are more fundamental, and don't seem to be solvable without significant changes to the policy enforced by the browser. In particular, there are several scenarios that indicate that a policy based on non-interference would have advantages over the current access control policy.

A first, very simple, motivating example for an information flow policy is a scenario where a website sends code to perform calculations on user private data.

Example 1 (Tax Calculator)**.** *Suppose the fictitious website* `http://taxcalc.com` *offers the service of pre-calculating the amount of tax one has to pay in function of income, age, marital status and so forth. The service sends an HTML form for entering the user's information, and a JavaScript program that will calculate the tax due based on the information entered in the form.*

The user wants assurance that the information he enters does not leave his computer — not even to the website providing the service. Obviously, the same-origin-policy does not offer any protection for this scenario. More fundamentally, if we assume that further interactions between the user and the website are essential (for instance to pay for the service), no access-control policy can provide this assurance: the script needs access to the private data to perform its function, and it needs access to the network to send invoicing information to the service. What is needed is an information flow enforcement mechanism that can ensure that the script can not leak private information to the network.

In many cases, the user will of course trust the website he is interacting with, and will be more concerned with information leaked to other sites.

Example 2 (Flight ticket)**.** *Consider an e-commerce site where users can order flight tickets. Obviously, the user will be fine with sharing some private information such as name, birth date and even credit-card information with the website. However, the user would like to have assurance that this information does not leak to other sites.*

The same-origin-policy provides some protection for this scenario: it will for instance ensure that scripts running in the user's browser and belonging to web pages from other origins can not access the information entered by the user. However, scripts that are part of the e-commerce web pages will have access and they can easily transmit information to other sites. This can be done by initiating an HTTP request to that other site where some information to be leaked is encoded in the URL or parameters of the request [9]. The script that leaks the information does not necessarily come from the trusted site. There are many ways in which malicious scripts can find their way into pages from trusted websites. Two common attack vectors are (1) cross-site scripting (XSS), where a vulnerability in the server software enables an attacker to inject scripts in the web pages served by the server [14], and (2) the inclusion of advertisements from third-party ad-providers; such advertisements are regularly implemented as scripts that run within the same origin as the including page [12].

Another example scenario is a combination of the two examples above: some information or resources that the web application user provides are private to the user, others are intended to be shared only with the web application provider.

Example 3 (Flight ticket (revisited))**.** *Even though the user trusts* `http://www.air.com` *with the information necessary to purchase a flight ticket, scripts from that site get access to other information that the user might want to protect, such as for instance geographical location (available to JavaScript through the geolocation API), or the clipboard contents. Hence the user has partial trust in the site and shares some information, but wants assurance that (1) the information shared with the site will not leak to other sites (as in the previous example), and in addition (2) some user private information accessible to scripts remains private to the user (as in the first example).*

An important additional challenge is the fact that for many web applications, some form of information flow between origins is actually desired. So any proposed browser security policy should not block such information flows. It is for instance common to include content (e.g. images and scripts) from other origins in web pages. A strict non-interference policy would prohibit such techniques and hence be strongly incompatible with the current web. It should for instance be possible for a script loaded from `b.com` into a page served from `a.com` to load images from any origin, since web advertising relies essentially on such scenarios. It should however be impossible for the script to leak information private to the user or to `a.com` in that scenario.

The examples above illustrate that non-interference is a promising candidate for a (baseline) browser security policy, but two important problems need to be addressed.

First, an enforcement mechanism for non-interference at the level of the browser is needed. While several browser security countermeasures based on information flow security techniques have been proposed, none of them can enforce non-interference for the full browser and for a broad class of security lattices in a secure and precise way (see the Related Work section for a detailed discussion). This paper proposes an enforcement mechanism, and proves it secure and precise.

Second, non-interference is parameterized with a policy: a poset of security levels, and an assignment of such levels to browser inputs and outputs. So an important problem is to select suitable policies. This paper analyzes several interesting policies and shows that they can securely handle the scenarios above, yet stay compatible with desired cross-origin information flows such as image and script loading.

3 Informal overview

To address the first problem (the development of a general, secure and precise enforcement mechanism for a full browser), we need a formal model of a browser. In order to experiment with policies, this browser model should be executable. *Featherweight Firefox* is a browser model developed by Bohannon and Pierce [2] that satisfies these two requirements. It is a small-step operational semantics of a browser, which is implemented in OCaml.[2]

In another paper, Bohannon et al. [3] defined several variants of non-interference suitable for browsers, and proposed a bisimulation-based proof technique to establish one of these types of non-interference (called *ID-security* in their paper). We will use their definition and proof technique to prove the security of our enforcement mechanism. Interested readers can find more information about the Featherweight Firefox and reactive non-interference in the appendix and in the original papers.

The enforcement mechanism we propose in this paper is based on a relatively new dynamic technique for achieving non-interference: secure multi-execution [4, 7]. The core idea of this mechanism is to execute the program multiple times (one copy of the program for each security level), and to ensure that (1) outputs of a given level l are only done in the execution at level l (outputs are suppressed in other copies), and (2) inputs at a level l are only done at level l (for the other copies above l, the values that were input by level l are reused, whereas copies that are not above l are fed a default input value). Hence the copy that does output at level l only sees inputs of levels below l and hence the output could not have been influenced by inputs of a higher level. Non-interference follows easily from this observation.

Devriese and Piessens [7] have worked out this enforcement mechanism for the case of a simple sequential programming language with synchronous I/O,

[2] We used the version available from Aaron Bohannon's webpage: `http://www.cis.upenn.edu/~bohannon/browser-model/` where interested readers can find the full definition of the model.

```
1 var a = parseInt(document.getElementById('a').value);
2 var b = parseInt(document.getElementById('b').value);
3 var sum = a + b;
4 document.getElementById('c').value = sum;
5 var url = 'http://attacker.com' + '?t=' + sum;
6 document.getElementById('banner').src = url;
```

Figure 1: JavaScript code example

and have proven security and precision in that setting. Capizzi et al. [4] have implemented the technique at the level of operating system processes for the case of two security levels.

The mechanism we propose adapts this technique to reactive systems, and we prove its security (somewhat weaker than what Devriese and Piessens have shown in their setting, because we lose termination- and timing-sensitivity), as well as its precision (somewhat stronger than the result by Devriese and Piessens, because we show precision under weaker assumptions).

Let us explain the mechanism by means of an example. Consider again the tax calculation example from Section 2. The JavaScript code in Figure 1 models the essence of this example: the user provides private inputs (two integers) in the text fields a and b , and the JavaScript code computes their sum and displays this in textfield c. We can assume this JavaScript code is part of an event handler that fires whenever the user changes the contents of a or b.

The code in the figure also shows a potential attack: the script will leak the (secret) sum to `http://attacker.com` by sending an HTTP request to that domain with the secret as a parameter (setting the `src` property of an image HTML element in JavaScript will have as a side effect that the image is reloaded from the URL assigned to the `src` property). Recall from Section 2 that the JavaScript code was not necessarily endorsed by the tax calculation site. It could have been injected through a cross-site scripting (XSS) attack or hidden in an advertisement running on the page.

Table 1 shows the behavior of Featherweight Firefox on this example. Since Featherweight Firefox does not support images, we simulate the information leak through a page load instead of an image load, which is from the point of view of information flow security the same thing.

If we assume that the inputs to the textfields have a high security level (H), and the output to `http://attacker.com` has a low security level (L), then this program is clearly not secure: high inputs leak to low outputs. How will our enforcement mechanism close this leak?

First, we have to assign security levels to all inputs and outputs of Featherweight Firefox; input to textfields is assigned H, and all the other inputs and outputs are L. Next, according to the idea of secure multi-execution, we run several copies of the web browser, one for each security level. Input events of level l are only processed by the copies with a level above or equal to l. Output events of level l are only produced in the copy at level l. Tables 2 and 3 show

Table 1: Correspondence between user's actions and I/O of Featherweight Firefox

Description of user actions and network events	Input/ Output of Featherweight Firefox
User opens a url of the tax calculator in a new window, as a result the new window is opened and an HTTP request is sent	`load_in_new_window`("`http://taxcalc.com`") ↳ `window_opened` ↳ `send`("taxcalc.com", *request_uri, cookies*, ""))
Network sends an HTTP response with html doc containing fields *a*, *b*, *c* and inlined JavaScript function	`receive`("taxcalc.com", 0, *cookie_updates*, *doc*(a=0, b=0, c=0, js_inline)) ↳ `page_loaded`(*user_window*, "`http://taxcalc.com`", *doc*(a=0, b=0, c=0, js_inline))
User types "2" into a text box *b*. This triggers the JavaScript event handler to execute the attack	`input_text`(*user_window*, 1, "2") ↳ `page_updated`(*user_window*, *doc*(a=0, b=2, c=2, js_inline)) ↳ `window_opened` ↳ `send`("attacker.com", *request_uri, cookies*, "?t=2")

what happens in both copies. The tables also show what inputs and outputs get suppressed in each level. For instance, for the L copy, the following things get suppressed: (1) the input events of level H (and hence also all output events that would have been the result of that input event), and (2) the output events at level H.

Table 2: Run of L copy of the browser.

	Input/Output
L H L	`load_in_new_window`("`http://taxcalc.com`") ↳ ~~`window_opened`~~ ↳ `send`("taxcalc.com", *request_uri, cookies*, ""))
L H	`receive`("taxcalc.com", 0, *cookie_updates*, *doc*(a=0, b=0, c=0, js_inline)) ↳ ~~`page_loaded`(*user_window*, "`http://taxcalc.com`", *doc*(a=0, b=0, c=0, js_inline))~~
H	~~`input_text`(*user_window*, 1, "2")~~
L L	(further low input) ↳ `send`("attacker.com", *request_uri, cookies*, "?t=0")

The offending output to `http://attacker.com` is suppressed, as the L copy never gets the input event where the user is typing secret data in the text box. In the table, we show that even if the script would try to send the contents of *a* and *b* later in the execution in response to further L input, the actual output

Table 3: Run of H copy of the browser.

	Input/Output
L	`load_in_new_window`("`http://taxcalc.com`")
H	`window_opened`
L	↳ ~~`send`("taxcalc.com", *request_uri, cookies*, "")~~
L	`receive`("taxcalc.com", 0, *cookie_updates*, *doc*(a=0, b=0, c=0, js_inline))
H	↳ `page_loaded`(*user_window*, "`http://taxcalc.com`", *doc*(a=0, b=0, c=0, js_inline))
H	`input_text`(*user_window*, 1, "2")
H	↳ `page_updated`(*user_window*, *doc*(a=0, b=2, c=2, js_inline))
H	↳ `window_opened`
L	↳ ~~`send`("attacker.com", *request_uri, cookies*, "?t=2")~~
L	(further L input)
L	↳ ~~`send`("attacker.com", *request_uri, cookies*, "?t=2")~~

sent to "`attacker.com`" would only contain the sum of the default values in both textfields. There is never any information flow from H inputs to L outputs.

4 Formalization

We propose to apply the approach of secure multi-execution to a reactive system first. Given an information-flow policy, we build a new reactive system that we call a *wrapper*. The wrapper internally runs multiple copies (*sub-executions*) of the original reactive system: one for each security level. When the wrapper consumes an input event, its security level is determined, and the input event is passed to those sub-executions that are allowed to see it, i.e. the sub-executions at a level higher than the input event's level. When a sub-execution at a security level produces an output, its security level is determined and only if the two levels are the same, the output is produced by the wrapper.

Throughout this section, we assume the reader is familiar with the notion of reactive system and ID-security [3]; the appendix recaps the definitions and theorems we use.

4.1 Secure multi-execution of reactive systems

The information-flow policy contains a partially ordered set of security levels $(\mathcal{L}, \leq)$ and a function $\mathrm{lbl} : \mathit{Act} \to \mathcal{L}$ assigning security levels to all inputs and outputs of the reactive system. The output $\cdot$ is an output invisible at all levels, and can be used to represent internal activity of the system. (For instance to return from a producer state to a consumer state without producing real output [3].)

A state of the wrapper is a tuple (R, L), where

$$\text{LOAD}\ \frac{R(l) \xrightarrow{i} P_l \qquad \text{if } \mathrm{lbl}(i) \leq l \text{ then } R'(l) = P_l \qquad \text{else } R'(l) = R(l) \text{ for all } l}{(R, \emptyset) \xrightarrow{i} (R', \mathit{Upper}(i))}$$

$$\text{OUT-P}\ \frac{R(l) \xrightarrow{o} P \qquad \mathrm{lbl}(o) = l}{(R, l :: L) \xrightarrow{o} (R[l \mapsto P], l :: L)} \qquad \text{OUT-C}\ \frac{R(l) \xrightarrow{o} C \qquad \mathrm{lbl}(o) = l}{(R, l :: L) \xrightarrow{o} (R[l \mapsto C], L)}$$

$$\text{DROP-P}\ \frac{R(l) \xrightarrow{o} P \qquad \mathrm{lbl}(o) \neq l}{(R, l :: L) \xrightarrow{\cdot} (R[l \mapsto P], l :: L)} \qquad \text{DROP-C}\ \frac{R(l) \xrightarrow{o} C \qquad \mathrm{lbl}(o) \neq l}{(R, l :: L) \xrightarrow{\cdot} (R[l \mapsto C], L)}$$

Figure 2: Semantics for secure multi-execution of a reactive system.

- R is a function mapping security levels to states of the reactive system, $R : \mathcal{L} \to \mathit{State}$. $R(l)$ is the state of the sub-execution at level l.
- L is the list of the levels of all the sub-executions that are in producer state (you can think of it as the scheduler's *ready queue*).

States $(R, \emptyset)$ are consumer states of the wrapper and states (R, L) with $L \neq \emptyset$ are producer states. The initial state of the wrapper is a state $(R, \emptyset)$ such that for all $l \in \mathcal{L}$, the state $R(l)$ is the initial state of the original reactive system.

Fig. 2 shows the semantics of the wrapper. When a new input event i arrives, it is passed to the copies at the levels in $\mathit{Upper}(i)$ (defined as the list of security levels higher than the level of a given input i), and the wrapper makes a transition to a producer state (rule [LOAD]). Once the wrapper is in producer state (R, L) it takes the first security level l from the list of levels L and gives the copy at this level a chance to proceed. If it produces an output at level l, it is also produced by the wrapper (rules [OUT-P] and [OUT-C]), otherwise a silent output $(\cdot)$ is produced instead (rules [DROP-P] and [DROP-C]). If the copy at the selected level l goes to a consumer state, then this level is removed from the L (rules [OUT-C] and [DROP-C]).

It is intuitively almost trivial to see why this construction guarantees non-interference. Output at any level l is only produced from the sub-execution at level l, which only gets to see input at level l or lower, so any leaks from input at higher levels is impossible. On the other hand, the sub-execution at a level l receives identical input on level l or lower as the original. Therefore, if the program is such that higher-level input does not influence lower-level output, then our construction will still produce the same output as the original. It is possible that the order of outputs will be reordered though. We will discuss both of these aspects (*security* and *precision*) in the next subsections.

4.2 Security

First, we show formally that our enforcement technique guarantees non-interference: for any reactive system and any information flow policy, the wrapper that we construct for it will never produce information leaks.

Theorem 4.1 (Security)**.** *All the states of the wrapper are ID-secure.*

All the proofs of the lemmas and theorems can be found in the Appendix.

With Bohannon et al.'s bisimulation-based proof technique [3, Theorem 4.5] (see also appendix), it suffices to prove that there exists an ID-bisimulation $\approx_l$ such that for every state of the wrapper (R, L), we have $(R, L) \approx_l (R, L)$. For a list of security levels L, we use the notation $L|_l$ to represent the list of levels l' in L such that $l' \leq l$.

Definition 4.1. *The state* (R_1, L_1) *is* l*-similar to the state* (R_2, L_2) *(written* $(R_1, L_1) \approx_l (R_2, L_2)$*) iff*

- $R_1 \approx_l R_2$ *meaning* $\forall l' \leq l : R_1(l') = R_2(l'))$*, and*
- $L_1|_l = L_2|_l$*.*

To prove that this is a bisimulation, we basically need to show that l-similar states produce outputs that are equal at level l or lower, and that the relation is maintained when they receive inputs that are equal up to level l.

Lemma 4.1. *This* l*-similarity relation is an ID-bisimulation.*

4.3 Precision

On the other hand, we need to prove that our enforcement mechanism is precise: since it will sometimes modify the behaviour of programs, we need to prove that it does this in a sensible way, i.e. it does not observably modify behaviour for programs that already are secure. We show precise formal results to explain exactly what we mean by this.

First, we need to define what we mean when saying that our enforcement mechanism *does not observably modify the behaviour of programs.* Important to notice is that even for well-behaved programs, the wrapper can change the relative order of output events at different security levels. We assume that any observer will only observe at a single security level. This assumption is valid for the policies we will consider in Section 5. Then, we define the observer-indistinguishable$_l$ relation that relates input or output streams that "look the same" for observers at security level l. Like Bohannon et al., we use a coinductive definition to clearly specify this definition for infinite streams.

Definition 4.2. *Define* observer-indistinguishable$_l(S, S')$ *coinductively with the*

following rules:

$$\text{observer-indistinguishable}_l([], [])$$

$$\frac{lbl(s) \neq l \qquad \text{observer-indistinguishable}_l(S, S')}{\text{observer-indistinguishable}_l(s :: S, S')}$$

$$\frac{lbl(s') \neq l \qquad \text{observer-indistinguishable}_l(S, S')}{\text{observer-indistinguishable}_l(S, s' :: S')}$$

$$\frac{\text{observer-indistinguishable}_l(S, S')}{\text{observer-indistinguishable}_l(s :: S, s :: S')}$$

This notion is weaker than Bohannon et al.'s ID-similarity. In fact, we have the following result:

Lemma 4.2. *If* $O \approx_l^{ID} O'$*, then* observer-indistinguishable$_{l'}(O, O')$ *for all* $l' \leq l$*.*

The notion of similarity that we will use for our precision results requires that the wrapper's output "looks the same" as the original output for observers at any one level. We define $S \approx_l^{obs} S'$ if observer-indistinguishable$_{l'}(S, S')$ for all $l' \leq l$ and $S \approx^{obs} S'$ if this holds for all l'. Note how these notions allow for changes in the relative order of events on different security levels.

Another notion we need is the projection of a finite stream at a certain security level l. The projection function π_l removes from the stream those events that are at a level not below l.

Definition 4.3. *Define, for finite* I_0

$$\pi_l([]) = []$$
$$\pi_l(i :: I_0) = \begin{cases} \pi_l(I_0) & \text{if } \text{lbl}(i) \not\leq l \\ i :: \pi_l(I_0) & \text{if } \text{lbl}(i) \leq l \end{cases}$$

Our enforcement mechanism produces observably equivalent outputs for those inputs for which the original reactive system is already "well-behaved" with respect to the security policy. We use the following precise definition:

Definition 4.4. *Given a reactive system state* Q *and a finite input* I *and output* O *such that* $Q(I) = O$*, we say that* Q *behaves securely for input* I *iff for all* $l \in \mathcal{L}$*, we have that* $Q(\pi_l(I)) = O_l$ *with* observer-indistinguishable$_l(O, O_l)$*.*

These are the definitions we need to state the first of our precision theorems. The following theorem is the most detailed result, and shows that for those inputs for which the reactive system behaves securely, the corresponding wrapper produces results that are observationally equivalent.

Theorem 4.2 (Precision for individual runs)**.** *Suppose a given reactive system state Q behaves securely for input I and $Q(I) = O_Q$. Define the corresponding wrapper $W = (R_Q, L)$ with $R_Q(l) = Q$ for all l, $L = \emptyset$ if $Q \in ConsumerState$ and $L = \mathcal{L}$ if $Q \in ProducerState$. For $O_W = W(I)$, we have that $O_Q \approx^{obs} O_W$.*

This theorem is actually not a typical precision result for an information flow enforcement technique, because it does not require non-interference of the original system, as would be more typical (see e.g. Devriese and Piessens [7]). Instead, the theorem gives a sufficient condition for an individual execution to "behave securely" and produce observationally equivalent results. However, we can show that the previous theorem is stronger, by showing that if the original system was non-interferent, then all of its executions "behave securely".

Lemma 4.3. *If a given reactive system state Q is ID-secure, then it behaves securely for any input I.*

This lemma easily leads to the following, more classical, precision theorem.

Theorem 4.3 (Precision)**.** *Suppose a given reactive system state Q is ID-secure, and $Q(I) = O$. Define the corresponding wrapper $W = (R_Q, L)$ with $R_Q(l) = Q$ for all l, $L = \emptyset$ if $Q \in ConsumerState$ and $L = \mathcal{L}$ if $Q \in ProducerState$. For $O' = W(I)$, we have that $O \approx^{obs} O'$.*

The stronger result is important in practice. Featherweight Firefox (without secure multi-execution) is never ID-secure: even if all scripts that have been loaded up to now behaved fine, somewhere in the future a malicious script might be loaded that leaks information. So the classical precision theorem does not apply, and it does not allow us to conclude precision for runs of the browser that actually behave well.

So what we need is a theorem that says: if the run of the browser up to some point behaved well, then our enforcement will not modify that run in an observable way. This is exactly what our first precision theorem does.

Note that we are only talking about precision here: security is never at stake. Featherweight Firefox with our enforcement mechanism will always be ID-secure. The point here is that we want to relate the behavior of the secured browser with the unsecured one, and this cannot be done with a classical precision theorem.

5 Information flow policies

The implementation of our information flow enforcement technique for Featherweight Firefox allows us to demonstrate some different information flow policies that are valuable as browser security policies. The three basic policies we show demonstrate on the one hand the power of information flow policies, allowing us to define precisely (contrary to traditional access control policies) the property that we want to enforce. On the other hand, our examples also show that it

Table 4: Simple High/Low policy

User input	`load_in_new_window`(*url*)	L
	`input_text`(*user_window, nat, string*)	H
User output	`window_opened`	H
	`page_loaded`(*user_window, url, rendered_doc*)	H
	`page_updated`(*user_window, rendered_doc*)	H
Network input	`receive`(*domain, nat, cookie_updates, resp_body*)	L
Network output	`send`(*domain, request_uri, cookies, string*)	L

is our enforcement technique that enforces the policies in such a way that non-complying programs are dealt with as precisely as possible (not just terminating them as traditional information flow policy enforcement techniques would).

The importance of current web applications calls for strong guarantees, but the enormous amount of legacy software out there calls for a highly compatible solution. We think that it is the combination of the accuracy of information flow policies together with the precise support for non-complying programs that makes a nice fit for the requirements of a browser security policy context.

5.1 High/Low Policy

Let us take another look at the tax calculator example. We would like the browser to guarantee to the user that his input does not leave his computer. However, it is our goal to do this without breaking the legacy website that uses existing DOM APIs and might interact with internet servers for downloading scripts and images.

Already for this simple scenario, an access control policy is not fine-grained enough to achieve these goals. An access control policy can prevent information leaks through server requests only by allowing requests to pass or not based on their content. Unlike an information flow policy, it cannot reason about what input the requests' content was constructed from and as such, it cannot distinguish requests that actually leak information from those that don't.

An information flow policy can be more fine-grained. In this case, the policy could classify all user input as secret information (H) and all network requests as public (L), disallowing URLs for dynamically added images to depend on secret information. Table 4 shows the classification of some I/O events. The policy then states that output events at level L (public) are not allowed to be influenced by input events at level H (secret).

Table 5 represents the execution of a prototypical script for the tax calculator example using our enforcement technique for the simple High/Low policy. In the example, the browser responds in the normal way to the user navigation and loads the page. In response to the user entering text in one of the text boxes, the script modifies the page and the browser shows the modified page to the

Table 5: Simple High/Low policy at work for a tax calculator script containing a malicious leak.

L	`load_in_new_window`("`http://taxcalc.com`")		
	L	L	`send`("taxcalc.com", *request_uri, cookies*, ""))
		H	~~`window_opened`~~
	H	L	~~`send`("taxcalc.com", *request_uri, cookies*, ""))~~
		H	`window_opened`
L	`receive`("taxcalc.com", 0, *cookie_updates*, *doc*(a=0, b=0, c=0, js_inline))		
	L	H	~~`page_loaded`(*user_window*, "`http://taxcalc.com`", *doc*(a=0, b=0, c=0, js_inline))~~
	H	H	`page_loaded`(*user_window*, "`http://taxcalc.com`", *doc*(a=0, b=0, c=0, js_inline))
H	`input_text`(*user_window*, 1, "2")		
	H	H	`page_updated`(*user_window*, *doc*(a=0, b=2, c=2, js_inline))
		H	`window_opened`
		L	~~`send`("attacker.com", *request_uri, cookies*, "?t=2")~~

user. However, the script is malicious and tries to leak the user's information by triggering a navigation to a page under the "`attacker.com`" domain.

Our enforcement mechanism modifies this behaviour: depending on the level of each input event (column 1), it is either passed to both or only one of the sub-executions at levels L and H. These sub-executions are shown under the input event with their security level in column 2. Each sub-execution produces output events (column 4 with the output's security level in column 3), but only the output events at the sub-execution's own level are actually produced by the wrapped system. This effectively prevents the information leak that the script is trying to trigger, as we can see, the request to "`attacker.com`" is not produced, since the low execution does not receive the `input_text` input event and the high execution will not trigger this output. To the people behind "`attacker.com`", it seems as though the user loads the page, but never inputs any text.

Even though the behaviour is being modified, the user in this case sees exactly the same behaviour as he would without the security mechanism. As long as future user-observable behaviour does not depend on the reactions of untrusted observers to information leaking requests, this is likely to be okay. For example in the case that the leak was the result of an XSS-attack [14], this is likely to be the case.

It might seem that our simple High/Low policy will block any request to a website. However, this is not the case. Intuitively, the reason that the request to "`attacker.com`" is being blocked is that it is being made in response to a user input event, and the fact that the user has performed a text input is defined to be private information by our policy. Toward observers on the L security level, the policy enforcement therefore replaces this behavior by default behavior coming from the L execution, which is kept under the illusion that no user input has occurred.

Table 6: Simple High/Low policy, third-party script

L	`load_in_new_window`("`http://taxcalc.com`")		
	L	H	~~`window_opened`~~
		L	`send`("taxcalc.com", *request_uri, cookies*, ""))
	H	H	`window_opened`
		L	~~`send`("taxcalc.com", *request_uri, cookies*, ""))~~
L	`receive`("taxcalc.com", 0, *cookie_updates*, *doc*(a=0, b=0, c=0, js_remote))		
	L	H	~~`page_loaded`(*user_window*, "`http://taxcalc.com`", *doc*(a=0, b=0, c=0, js_remote))~~
		L	`send`("remote.com", *request_uri, cookies*, ""))
	H	H	`page_loaded`(*user_window*, "`http://taxcalc.com`", *doc*(a=0, b=0, c=0, js_remote))
		L	~~`send`("remote.com", *request_uri, cookies*, ""))~~
L	`receive`("remote.com", 0, *cookie_updates*, js_function)		
	L H		
H	`input_text`(*user_window*, 1, "2")		
	H	H	`page_updated`(*user_window*, *doc*(a=0, b=2, c=2, js_remote))
	H	H	`window_opened`
		L	~~`send`("attacker.com", *request_uri, cookies*, "?t=2")~~

We illustrate this observation by demonstrating another script that makes legitimate use of external requests. Let's suppose that the tax calculator script needs information from a third-party website (for example, an up to date table of tax rates for different income ranges). In Table 6, we show what happens if the script requests to download such a table from a server at "`remote.com`" after the page has loaded.

In this case, we see that the legitimate request to "`remote.com`" is allowed to pass without any problem. When its response comes in, both the L and H execution see this and the program behaves as intended. This shows that indeed, we get what we ask for: our policy specifies only that user input must not influence network output, and indeed we see that network output that is not influenced by user input behaves as intended.

Finally, the simple High/Low policy is designed specifically for a case like a tax calculator website where all interaction with the user can be assumed to occur inside the browser. Specifically, the policy assumes all information coming from the network as public data. It is important to remember, when applying this policy, that this policy will not provide any protection for information that is classified as public, like for example page content or cookies. As an example, Table 7 shows a case where a cookie tracking the user's language is leaked. This is not a limitation of our enforcement mechanism, but a limitation of the specific policy that it is enforcing here.

Note that this policy is a very simple information flow policy, but already

Table 7: Simple High/Low policy, Stealing cookies

...			
L	`receive`("taxcalc.com", 0, ("lang", "ru"), *doc*(a=0, b=0, c=0, js_remote))		
	L	H	~~`page_loaded`(*user_window*, "http://taxcalc.com", *doc*(a=0, b=0, c=0, js_remote))~~
		L	`send`("remote.com", *request_uri, cookies*, ""))
		L	`send`("attacker.com", *request_uri, cookies*, "?lang=ru"))
	H	H	`page_loaded`(*user_window*, "http://taxcalc.com", *doc*(a=0, b=0, c=0, js_remote))
		L	~~`send`("remote.com", *request_uri, cookies*, ""))~~
		L	~~`send`("attacker.com", *request_uri, cookies*, "?lang=ru"))~~
...			

achieves something that was previously not possible with simple access control policies. We can run a website making sure that certain user information is never leaked. It is not hard to think of interesting extensions of this idea for real-life browser scenario's: for example a "Keep all information in this field inside my browser" button that you can push to prevent information entered into a field from leaving your browser. The browser's policy enforcement could then use an enforcement technique like ours to guarantee security of the information, and in many cases without affecting the further behaviour of the site.

5.2 Separating origins

The airplane tickets e-commerce site example is more typical for a general web site. In this scenario, a level of trust is assumed between the user and the company hosting the ticketing website, in order for the ticketing company to provide useful information or services. Nevertheless, the standard same-origin-policy (SOP) is not sufficient as it allows (in practice) this data to be sent anywhere.

We believe that the basic model of the SOP is actually correct. When a user enters information on a website, it is typically his intent to disclose this information to the owner of that website, but not others. Likewise, information received from a website can be trusted to be sent back to this website but not to others. We think that using information flow enforcement techniques such as the one described in this paper, a replacement for the SOP can be defined that does achieve the intended information protection, with sufficient backwards compatibility for much of the code currently "out there".

A somewhat evident idea here is to use a security lattice with three types of levels: L, M(dom) for any domain *dom* and H. The L and H levels are smaller resp. bigger than all others and the M(...) domains are mutually incomparable. The M(dom) level is assigned to all network events originating from or going to this domain and to all user input events that contain information destined for a

Table 8: Origin separation policy

User input	`load_in_new_window`(*url*)	L
	`input_text`(*user_window, nat, string*)	M(...)
User output	`window_opened`	H
	`page_loaded`(*user_window, url, rendered_doc*)	H
	`page_updated`(*user_window, rendered_doc*)	H
Network input	`receive`(*dom, nat, cookie_updates, resp_body*)	M(dom)
Network output	`send`(*dom, request_uri, cookies, string*)	M(dom)

page on this domain. Output events going to the user are classified as H. This policy is summarized in Table 8.

Table 9 shows the execution of a prototypical airline ticketing website script under the origin separation policy from Table 8. We see that network output to "`air.com`" is now permitted to be influenced by information from user input in the corresponding web page.

Something interesting happens when we consider a page that tries to download a third-party script at page load time. In Table 10, we see that our security mechanism prevents the request for the third party script from being sent at all, very likely breaking the site's behaviour.

Let us see why this happens. The request for the third-party script on host "`remote.com`" is produced in response to the `receive` input event representing the receival of the website document. However, our policy marks this input event as information that must only be revealed to the "`air.com`" domain. Hence the request to the "`remote.com`" should not be sent. Our security enforcement cannot be blamed for breaking the website, since it is only executing what the policy specified. So that must mean that the policy is wrong and we should have classified the `receive` input event on a different security level?

Unfortunately, there is a very good reason why it should be classified at this level. If we suppose the page that is received represents the third step in the airline ticket purchasing process, and contains a summary of all data previously input by the user, then this is clearly information that we want to protect and the policy is correct to not just allow this info to leak to third-party sites.

One solution would be to to declassify parts of the page. However, declassification is complex [16] and makes it hard to understand the policy that is still being enforced. So we prefer to avoid it.

A possible answer is that the information flow policy is not fine-grained enough. If we want to refine the SOP retaining maximum compatibility, we need to define a policy that does a better job of formalizing the assumptions in the current web security model. In this case, there is the implicit notion that an HTML document contains information at different confidentiality levels. If the document specifies that it requires a certain script to function then this information must be permitted to leak to the website in question. However, we

Table 9: Origin separation policy. M1 = M("air.com"), M2 = M("attacker.com").

L	`load_in_new_window`("`http://air.com`")		
	L	H	~~`window_opened`~~
		M1	`send`("air.com", *request_uri, cookies*, ""))
	H	H	`window_opened`
		M1	~~`send`("air.com", *request_uri, cookies*, ""))~~
M1	`receive`("air.com", 0, *cookie_updates*, *doc*(age=0, ..., js_inline))		
	M1	H	~~`page_loaded`(*user_window*, "`http://air.com`", *doc*(age=0, ..., js_inline))~~
	H	H	`page_loaded`(*user_window*, "`http://air.com`", *doc*(age=0, ..., js_inline))
M1	`input_text`(*user_window*, 0, "25")		
	M1	H	~~`page_updated`(*user_window*, *doc*(age=25, ..., js_inline))~~
		H	~~`window_opened`~~
		M1	`send`("air.com", *request_uri, cookies*, "?t=25")
		H	~~`window_opened`~~
		M2	~~`send`("attacker.com", *request_uri, cookies*, "?t=25")~~
	H	H	`page_updated`(*user_window*, *doc*(age=25, ..., js_inline))
		H	`window_opened`
		M1	~~`send`("air.com", *request_uri, cookies*, "?t=25")~~
		H	`window_opened`
		M2	~~`send`("attacker.com", *request_uri, cookies*, "?t=25")~~

Table 10: Origin separation policy, third-party script. M1=M("air.com"), M2=M("remote.com").

L	`load_in_new_window`("`http://air.com`")		
	L	H	~~`window_opened`~~
		M1	`send`("air.com", *request_uri, cookies*, ""))
	H	H	`window_opened`
		M1	~~`send`("air.com", *request_uri, cookies*, ""))~~
M1	`receive`("air.com", 0, *cookie_updates*, *doc*(age=0, ..., js_remote))		
	M1	H	~~`page_loaded`(*user_window*, "`http://air.com`", *doc*(age=0, ..., js_remote))~~
		M2	~~`send`("remote.com", *request_uri, cookies*, ""))~~
	H	H	`page_loaded`(*user_window*, "`http://air.com`", *doc*(age=0, ..., js_remote))
		M2	~~`send`("remote.com", *request_uri, cookies*, ""))~~
...			

have to ensure that the rest of the document cannot be leaked in the process. The next subsection discusses how to incorporate this in our mechanism.

5.3 Sub-input-event security policies

The key to solving the issue is to be able to assign different labels to different parts of a single input event. One simple solution is to model such an input event as a number of separate input events, so that we can give each of these parts a different level. Then our enforcement mechanism and our security and precision theorems can be applied as before.

An alternative, more intuitive way of thinking about this splitting of an input event (where different levels can see a different subset of the parts of the splitted event), is to consider security-level dependent projections that project an input event on the part of the event visible to a specific level. We discuss our solution from this angle.

So far, the information flow policy was defined by the function $\text{lbl} : \mathit{Act} \to \mathcal{L}$ returning the security level for any input or output event.

Alternatively, we can give a more flexible definition of a security policy by restricting the lbl function to output events and additionally requiring a set of projection functions $\{\pi_l : \mathit{Input} \to \mathit{Input} \cup \{\mathit{Suppress}\} \mid l \in \mathcal{L}\}$. These projection functions have to be idempotent and such that the projection function π_H at maximum level H is the identity function and for all $l \leq l'$ and input events i and i', we have that $\pi_l(i) = \pi_l(i')$ whenever $\pi_{l'}(i) = \pi_{l'}(i')$.

The projection function π_l intuitively defines the view of an input event at security level l ($\mathit{Suppress}$ is a keyword signaling that the input event should not be visible at this level) and what we've been doing up to now actually corresponds to projection functions defined as follows:

$$\pi_l(i) = \begin{cases} i & \text{if } \text{lbl}(i) \leq l \\ \mathit{Suppress} & \text{if } \text{lbl}(i) \not\leq l \end{cases}$$

Based on these projection functions, we can adapt the formal definitions of non-interference and our enforcement technique in an obvious way. The wrapper constructed by our enforcement technique will no longer send an input event i to all sub-executions at levels $l \geq \text{lbl}(i)$, but will instead send the projection $\pi_l(i)$ of the input event to all sub-executions at levels l such that $\pi_l(i) \neq \mathit{Suppress}$.

With these sub-input-event policies, we can refine our solution for the origin separation policy to make it align more closely with the assumptions that are implicit in the model of a web browser. In particular, for an input event $i = \texttt{receive}(\mathit{dom}, \mathit{nat}, \mathit{cookie\_updates}, \mathit{body})$ where body is an HTML document, we

propose to define the projection functions as follows:

$$\pi_l(i) = \begin{cases} i & \text{if } l = H \text{ or } l = M(dom), \\ \texttt{receive}(dom, nat, \{\}, proj_{dom'}(body)) & \\ & \text{if } l = M(dom'), dom \neq dom', \\ Suppress & \text{if } l = L, \end{cases}$$

where $proj_{dom'}(body)$ projects an HTML document onto an almost empty document with only public information remaining. It is important that script tags referencing scripts on domain dom' are also kept.

The assumptions here is that when a server generates a page for a user, most information is intended to be kept private, where private means that it can only flow to the user and back to the originating server. We make an exception only for that information that the server must have intended to be leaked to certain destinations. For example, when the server links to a script on a third-party domain, then the server is well aware that this will trigger a request to the domain in question. Therefore, there is no harm in disclosing this fact on the security level corresponding to that third-party domain.

For incoming scripts, we propose for now that for input events of the form $i = \texttt{receive}(dom, nat, cookie\_updates, body)$ where $body$ is a script, we model these as public input, i.e. $\pi_l = i$ for all l. This models the assumption that loaded scripts do not contain private information and as such, their content does not require protection. This assumption is valid for many typical scripts like public JavaScript libraries and corresponds to the behaviour of the current browser security policies, making the policy backwards-compatible.

If we apply this relaxed policy to the "`air.com`" website script with a third-party script, then contrary to Table 10, the request to "`remote.com`" will be sent properly, in the M2 execution. Table 11 further shows how the response for this event will be classified at security level M2, and the behaviour of the website will be as intended, but document data is well protected.

Scenario's exist where loaded scripts do contain private information (we think of, for example, JSON replies from web services). In such cases, protection of the returned script is valuable, since there might be a form of CSRF "JavaScript hijacking" [6] attack taking place. Taking inspiration from existing heuristics to protect client-side against CSRF attacks [6], we believe an interesting path for future work is to investigate heuristics classifying the contents of loaded third-party scripts in a reasonably secure, yet compatible way. Such heuristics can take into account information like the security properties of the protocol the script was downloaded over (e.g. HTTPS vs. HTTP), whether cookies were sent along with the request, etc. In some recent standards, such as Cross-Origin Resource Sharing (CORS), the server will even indicate cross-domain accessibility of the http response content. If cookies are available to be sent along with the request, it might even be a good idea to request the script twice (resp. with and without cookies) and expose the two results on different

Table 11: Fine-grained origin separation policy, third-party script. M1=M("air.com"), M2=M("remote.com").

L	`load_in_new_window`("`http://air.com`")		
	L	H	~~`window_opened`~~
		M1	`send`("air.com", *request_uri, cookies*, ""))
	H	H	`window_opened`
		M1	~~`send`("air.com", *request_uri, cookies*, ""))~~
M1	`receive`("air.com", 0, *cookie_updates*, *doc*(age=0, ..., js_remote))		
	M1	H	~~`page_loaded`(*user_window*, `"http://air.com"`, *doc*(age=0, ..., js_remote))~~
		M2	~~`send`("remote.com", *request_uri, cookies*, ""))~~
	M2	H	~~`page_loaded`(*user_window*, `"http://air.com"`, $proj_{M2}$(doc(age=0, ..., js_remote)))~~
		M2	`send`("remote.com", *request_uri,* {}, "")
	H	H	`page_loaded`(*user_window*, `"http://air.com"`, *doc*(age=0, ..., js_remote))
		M2	~~`send`("remote.com", *request_uri, cookies*, ""))~~
M2	`receive`("remote.com", 0, {}, js_function)		
	...		

levels (technically, this goes even further than the projection functions model). Experiments with real websites are needed to flesh this out further, and we leave further refinement of this policy for future work.

6 Implementation

We have implemented our enforcement mechanism based on the idea of secure multi-execution for the Featherweight Firefox browser model in OCaml[3]. We tested it on our policies described in Section 5 using the examples from the same section.

The operational semantics defined in Fig. 2 is implemented at the level of a reactive system. Even though the construction is relatively simple, we had to address several interesting issues.

One problem is the fact that policies can potentially have an infinite number of levels, for instance in the policies where there is one level per origin. We have solved this by adding new levels lazily. The following lemma shows that this is sound (we prove it for finite but unbounded streams).

Recall that the state of the wrapped system is a tuple (R, L), where R is a pointer to the running copies of the browser and $L \neq \emptyset$ is a list of levels at which the browser copies are in producer state. When waiting for a new input,

[3]It can be accessed here: `http://disi.unitn.it/~bielova/sme-firefox`.

the state of the wrapped system is $(R, \emptyset)$ because all the current states of the browser copies are in consumer state.

Lemma 6.1. *Suppose the current state of the wrapped system is $W = (R, \emptyset)$. If R is defined for l and l' and the input I is such that $\pi_l(I) = \pi_{l'}(I)$ (the input I looks the same at l and l'), then $R(l) = R(l')$.*

Hence, when the browser first sees an input of level l (e.g. the user visits a new website), we can lazily add the sub-execution at that level based on this lemma. We take the existing sub-execution at level l' such that $\pi_l(I) = \pi_{l'}(I)$ (where I is the input arrived so far) and use a copy of it for l. In the implementation it would usually mean that l' is one level lower than l in the security lattice.

In particular, we implemented this idea for the policy with security levels of three types: L, M(dom) and H. In the beginning of the run the wrapped system has only two copies of the original browser: at levels H and L. If the input at a new level l arrives, then in the security lattice it is such that $\mathrm{L} < l < \mathrm{H}$. This means that all the inputs I arrived so far look the same at new level l and the lowest level L: $\pi_l(I) = \pi_L(I)$. Hence, we clone the L copy of the browser and use it for a new sub-execution at the new level l.

Do copies at level l' always exists for an input at a new level l such that $\pi_l(I) = \pi_{l'}(I)$? Unfortunately not. If the new input level l is such that there exist two incomparable levels $l_1 < l$ and $l_2 < l$, then $\pi_l(I) \neq \pi_{l_1}(I)$, $\pi_l(I) \neq \pi_{l_2}(I)$. In this case the lth copy of the browser should be a combination of the l_1th and l_2th copies. Note that this cannot happen if we consider security level lattices (instead of general posets) and if we ensure that the set of levels for which we run a copy remains a sub-lattice after adding a new level.

Our implementation only considers the security posets of the forms we discussed in the previous sections, and for these posets, we can always lazily add new levels. This is because the security level posets we consider are in fact lattices and any set containing H and L is a sub-lattice.

Similarly, it can happen that the output event o produced by the lth copy of the browser has a level l' and is not in the list of levels at which the wrapped system is currently running the browser copies. This means that there is no copy of the browser at level l' that will be able to output o. Hence, if the level l' of o is such that $l < l'$ then we allow the lth copy to output o because this does not violate the non-interference (higher inputs should not influence lower outputs). However, we do not add a new copy of the browser at level l' to be run by the wrapped system because it is not necessary for the correctness of our approach.

Note that in our implementation, the enforcement mechanism is applied to the full browser as the underlying reactive system. For a real implementation, this is technically possible (note the similarity to Capizzi et al.'s shadow executions [4]), but we think there may be technical advantages to applying the enforcement at a lower level, taking a single script as the system to be wrapped and performing the enforcement inside the browser. An advantage of the cur-

rent model is that no trust is placed in the implementation of the browser and that it is compatible with the design of Featherweight Firefox.

7 Related Work

There is a large body of related work on information flow security in general, or on web security techniques in general. We refer the reader to three good sources where these fields are surveyed. Sabelfeld and Myers [15] survey static techniques for information flow enforcement, and Le Guernic [8] surveys dynamic techniques. The PhD thesis of Martin Johns [10] gives a good survey of web security techniques and countermeasures for web-related vulnerabilites.

In some recent works, not yet covered in the surveys cited above, authors have developed dynamic [1] or static [5] techniques to enforce information flow security in a browser context. These techniques lack the precision guarantees of secure multi-execution, but on the other hand, secure multi-execution is likely to have a higher performance penalty.

In the rest of this section, we focus on the work that is most closely related to ours.

A first very related line of work is the work by Bohannon et al. which has been discussed extensively in appendix. Next, there are several other security countermeasures that have strong similarities to our approach.

The technique of secure multi-execution proposed by Devriese and Piessens [7] is the most closely related. These authors proposed secure multi-execution for enforcing noninterference, and proved it to be sound and precise. They show these results for a simple sequential programming language with synchronous I/O. Our work extends theirs, by showing how the same technique can be applied to reactive systems and hence browsers. Interestingly, the formal guarantees we get are different. Whereas Devriese and Piessens can prove timing-sensitive non-interference, we have to settle for termination-insensitive non-interference. The main reason for this is that we are more restricted in the reordering of output events. On the other hand, we get a substantially stronger precision result. We show precision for any well-behaved run, whereas Devriese and Piessens can only prove precision for programs that are termination-sensitively non-interferent.

A similar approach was proposed by Capizzi et al. [4] where they run two executions of operating system processes for the H (secret) and L (public) security level. They limit themselves to this simple two-element poset, but they provide an actual implementation, and report on benchmarks.

Some other web-browser security techniques solve different problems, but use techniques that look like ours.

One recently proposed technique called AdJail [12] is particularly aimed at the information flow between user data displayed on the web page and third-party advertisements. Similarly to the secure multi-execution and to the shadow execution approaches, the authors propose to have a shadow copy of the web page where all the interactions between the ad script and the original page are controlled. They also report on an implementation. This paper is an excellent

example of how shadow executions can be used to address specific web security issues at a reasonable performance cost, but obviously the scope of the protection offered is smaller than for our proposed countermeasure.

Doppelganger [17] is another example of a similar approach with a very specific focus. It focusses on keeping control over HTTP cookies. It suggests to have two copies of the running web page: with and without cookies. It defines the difference between the pages in such a way that in certain cases the enforcement mechanism can decide automatically whether keeping cookies is important for the correct functionality of the web page. This technique, however, is concerned only about HTTP cookies and does not cover general information flow or other browser functionalities.

Several very recent works have applied information flow analysis to web mashups. Magazinius et al. [13] propose an approach to construct a security lattice for mashups. Similarly to our approach, where an element of the security lattice depends on the origin of the event, the authors of this paper defined the elements as sets of origins. So the security lattice consists of elements with one origin (for events) and elements with all possible combinations of the origins. In cases where different origin domains have to communicate, the approach relies on declassification. The paper is focused on the definition of the policies, and does not focus on enforcement mechanisms.

Li, Zhang and Wang also deal with mashups in their Mash-IF approach [11]. The security levels there consists of a tuple of sensitivity level and an origin. It is a practical approach, but no soundness or precision guarantees are provided.

8 Conclusion and future work

This paper has studied the suitability of non-interference as a replacement for the same-origin-policy in browsers. We have shown that it is possible to enforce non-interference for a browser securely and precisely for a broad class of information flow policies (even including policies with an infinite number of levels). In addition we have shown that, even without any support for declassification, useful information flow policies for a browser can be defined.

An important remaining challenge is the development of efficient implementation techniques for our enforcement mechanism (or alternative secure and precise mechanisms). Another important item for future work is the evaluation of the impact of the policies we proposed on real web sites: while the security benefits of a non-interference policy are high, it is to be expected that there will be a price to pay. Even though we have shown by example that some level of compatibility with the current web can be maintained, it is to be expected that many detailed incompatibilities will show up, and evaluating the cost of these – and how they could be mitigated – is a key challenge for future work.

Still, we do believe that the results reported in this paper provide evidence that it is worthwhile to go further down this road, and do the substantial effort of integrating non-interference mechanisms in standard browsers in order to evaluate performance and compatibility costs.

Acknowledgment

This work has been partly supported by the EU under the projects EU-IP-MASTER, EU-FET-IP-SecureChange and EU-NoE-NESSoS. It is also partially funded by the Interuniversity Attraction Poles Programme Belgian State, Belgian Science Policy, and by the Research Fund K.U.Leuven. Dominique Devriese holds a Ph. D. fellowship of the Research Foundation - Flanders (FWO).

References

[1] Thomas H. Austin and Cormac Flanagan. Permissive dynamic information flow analysis. In Proceedings of the 5th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security, PLAS '10, pages 3:1–3:12, New York, NY, USA, 2010. ACM.

[2] A. Bohannon and B. C. Pierce. Featherweight firefox: Formalizing the core of a web browser. In Proceedings of the USENIX Conference on Web Application Development 2010, 2010. To be published.

[3] A. Bohannon, B. C. Pierce, V. Sjöberg, S. Weirich, and S. Zdancewic. Reactive noninterference. In Proceedings of the 16th ACM Conference on Communications and Computer Security, pages 79–90. ACM Press, 2009.

[4] R. Capizzi, A. Longo, V. N. Venkatakrishnan, and A. Prasad Sistla. Preventing information leaks through shadow executions. In Proceedings of 24th Annual Computer Security Applications Conference, ACSAC '08, pages 322–331. IEEE Computer Society, 2008.

[5] Ravi Chugh, Jeffrey A. Meister, Ranjit Jhala, and Sorin Lerner. Staged information flow for javascript. In Proceedings of the ACM SIGPLAN 2009 Conference on Programming Language Design and Implementation, volume 44, pages 50–62. ACM Press, 2009.

[6] Philippe De Ryck, Lieven Desmet, Thomas Heyman, Frank Piessens, and Wouter Joosen. Csfire: Transparent client-side mitigation of malicious cross-domain requests. In Lecture Notes in Computer Science, volume 5965, pages 18–34. Springer Berlin / Heidelberg, February 2010.

[7] D. Devriese and F. Piessens. Non-interference through secure multi-execution. In Proceedings of the 2010 IEEE Symposium on Security and Privacy, SP '10, pages 109–124. IEEE Computer Society Press, 2010.

[8] G. Le Guernic. Confidentiality Enforcement Using Dynamic Information Flow Analyses. PhD thesis, Kansas State University, 2007.

[9] Martin Johns. On javascript malware and related threats. Journal in Computer Virology, 4:161–178, 2008.

[10] Martin Johns. Code Injection Vulnerabilities in Web Applications - Exemplified at Cross-site Scripting. PhD thesis, University of Passau, 2009.

[11] Z. Li, K. Zhang, and X. Wang. Mash-IF : Practical Information-Flow Control within Client-side Mashups. In Proceeedings of IEEE/IFIP International Conference on Dependable Systems and Networks (DSN-10), pages 251–260, 2010.

[12] M. T. Louw, K. T. Ganesh, and V. N. Venkatakrishnan. AdJail : Practical Enforcement of Confidentiality and Integrity Policies on Web Advertisements. In Proceedings of the 19th USENIX Security Symposium, 2010.

[13] J. Magazinius, A. Askarov, and A. Sabelfeld. A lattice-based approach to mashup security. In Proceedings of ACM Symposium on Information, Computer and Communications Security (ASIACCS-10), pages 15–23. ACM Press, 2010.

[14] F. Nentwich, N. Jovanovic, E. Kirda, C. Kruegel, and G. Vigna. Cross-site scripting prevention with dynamic data tainting and static analysis. In Proceedings of the Symposium on Network and Distributed System Security (NDSS 2007), 2007.

[15] A. Sabelfeld and A. C. Myers. Language-based information-flow security. In IEEE Journal on Selected Areas in Communication, volume 21, pages 5–19, 2003.

[16] Andrei Sabelfeld and David Sands. Declassification: Dimensions and principles. J. Comput. Secur., 17:517–548, October 2009.

[17] U. Shankar and C. Karlof. Doppelganger: Better browser privacy without the bother. In Proceedings of the 13th ACM Conference on Communications and Computer Security, CCS '06, pages 154–167. ACM Press, 2006.

[18] K. Singh, A. Moshchuk, H.J. Wang, and W. Lee. On the incoherencies in web browser access control policies. In Proceedings of the 2010 IEEE Symposium on Security and Privacy, pages 463–478. IEEE Computer Society Press, 2010.

Appendix

8.1 Background

This appendix summarizes the work by Bohannan et al. on Featherweight Firefox [2] and reactive non-interference [3] that we build on in this paper, and contains no original contributions.

8.2 Reactive systems

At the highest level of abstraction, a browser is modeled as a *reactive system*, a particular kind of automaton that reacts to inputs by changing state and emitting outputs.

Definition 8.1. *A reactive system is a tuple*

$$(ConsumerState, ProducerState, Input, Output, \rightarrow)$$

where $\rightarrow$ is a labelled transition system whose states are $State = ConsumerState \cup ProducerState$ and whose labels are $Act = Input \cup Ouput$, subject to the following constraints:

- *for all $C \in ConsumerState$, if $C \xrightarrow{a} Q$, then $a \in Input$ and $Q \in ProducerState$,*
- *for all $P \in ProducerState$, if $P \xrightarrow{a} Q$, then $a \in Output$,*
- *for all $C \in ConsumerState$ and $i \in Input$, there exists a $P \in ProducerState$ such that $C \xrightarrow{i} P$, and*
- *for all $P \in ProducerState$, there exists an $o \in Output$ and $Q \in State$ such that $P \xrightarrow{o} Q$.*

Consumer states are the states where the system is idle and waiting for inputs. A reactive system can only handle one input event at a time (thus correctly modeling the fact that JavaScript event handlers are single threaded). Producer states are states where the system is processing, and from such producer states, the system can emit outputs. The definition allows for non-termination: it is possible that the system never returns to a consumer state.

Reactive systems transform streams of input events into streams of output events in the obvious way. A *stream* is defined as a coinductive interpretation of the grammar

$$S ::= [] \mid s :: S \tag{1}$$

where s ranges over stream elements. So a stream is a finite or infinite list of elements. We use metavariables I and O to range over streams of inputs i and outputs o, respectively. The *behavior* of a reactive system in a state Q is defined as a relation between the input streams and output streams.

Definition 8.2. *Coinductively define $Q(I) \Rightarrow O$ (Q transforms the input stream I to the output stream O) by the following rules:*

$$C([]) \Rightarrow [] \qquad \frac{C \xrightarrow{i} P \qquad P(I) \Rightarrow O}{C(i :: I) \Rightarrow O} \qquad \frac{P \xrightarrow{o} Q \qquad Q(I) \Rightarrow O}{P(I) \Rightarrow o :: O}$$

8.3 Featherweight Firefox

The notion of reactive system is very abstract. To analyze potential security policies, we should define a browser model that concretizes the abstract states, inputs and outputs. The *Featherweight Firefox* browser model [2] does exactly that. It includes many browser features such as multiple browser windows; cookies; sending HTTP requests and receiving HTTP responses; essential HTML elements such as text boxes, buttons and links; building document node trees (i.e. a simple variant of the Document Object Model), and also the basic features of JavaScript.

Featherweight Firefox (FF) is a reactive system, with a much more detailed definition of the input and output events, and the internal state of the browser. To understand our security proof further in the paper, it suffices to understand FF at the abstraction level of reactive systems. However, to understand the example policies and scenarios, some basic understanding of the full FF model is needed. We focus here on explaining the I/O events modeled by FF, the details of the internal browser state are less relevant.

When the browser is in a ConsumerState, it is ready to process any input event that could arrive. Input events can either come from the user (loading a URL in a new window, entering text in a text box, clicking a button), or from the network (receiving an HTTP response). Output events can also go to the user (web page is rendered or updated, window is closed) or to the network (sending HTTP request). The FF browser model defines precisely how the browser will react to each of these inputs by emitting outputs.

Some selected input and output events are shown in Table 12. We simplify some syntactical constructs to give the reader a better understanding without unnecessary details.

Table 12: Selected user and network I/O events.

User input	`load_in_new_window`(*url*)
	`input_text`(*user_window, nat, string*)
User output	`window_opened`
	`page_loaded`(*user_window, url, rendered_doc*)
	`page_updated`(*user_window, rendered_doc*)
Network input	`receive`(*domain, nat, cookie_updates, resp_body*)
Network output	`send`(*domain, request_uri, cookies, string*)

An input event `load_in_new_window` models the case where a user navigates to some URL. When a user inputs some text into a text box this is modeled by the event `input_text`. The second parameter is an index that uniquely identifies the text box that receives the input, and the third parameter is the text that was typed.

The display of a new window (`window_opened`) is an output event to the user and has no parameters since new windows are always created with a URL `about:blank`. Then after an html document is loaded or updated there may be

visible changes of the rendered document on the screen. The `page_loaded` and `page_updated` events model these outputs to the user; the *rendered_doc* parameter of these events models a rendered document, and contains only elements that are visible to the user (for instance, script source code is not visible in a rendered html document).

Sending an HTTP request is modeled as an output event `send` of the browser. The request is sent to a particular *domain* and the tuple (*request_uri, cookies, string*) models a simplified version of an HTTP request.

The only input that can come from the network is the reception of an HTTP response. It contains a domain, an index that uniquely defines the open network connection on which the response arrives, updated cookies and the actual response content (either an html document or a script).

The FF model is surprisingly rich. We will see examples including for instance the execution of event handlers implemented as scripts in an html page, further in the paper.

8.4 ID-security, or reactive non-interference

It remains to define what it means for a reactive system (and hence FF) to be non-interferent. Bohannon et al. investigate different notions of non-interference (for instance a termination sensitive and a termination insensitive notion), and use a notation that can distinguish the different notions. This paper only uses their notion of *ID-security*, a termination insensitive variant of non-interference. We specialize their definitions and notation to this case.

Let us assume that a poset of security levels is given. The predicate $visible_l(s)$ models what observers of security level l can see: $visible_l(s)$ is true iff the stream element s is visible to an observer at level l. For instance, if we think of the input stream elements as arriving from input channels with a given security level, then $visible_l(s)$ is true if l is higher than the level of s's channel in the security ordering.

First, we define what it means for two (input or output) streams to be equivalent up to security level l.

Definition 8.3. *Coinductively define $S \approx_l^{ID} S'$ (S is ID-similar to S' at l) with the following rules:*

$$\frac{}{[] \approx_l^{ID} []} \qquad \frac{visible_l(s) \qquad S \approx_l^{ID} S'}{s :: S \approx_l^{ID} s :: S'} \qquad \frac{\neg visible_l(s) \qquad S \approx_l^{ID} S'}{s :: S \approx_l^{ID} S'}$$

$$\frac{\neg visible_l(s) \qquad S \approx_l^{ID} S'}{S \approx_l^{ID} s :: S'}$$

Then, we can define when a reactive system (in a specific state Q) is secure.

Definition 8.4. *A state Q is ID-secure or (reactive) non-interferent if, for all l, $I \approx_l^{ID} I'$ implies $O \approx_l^{ID} O'$ whenever $Q(I) \Rightarrow O$ and $Q(I') \Rightarrow O'$.*

As Bohannon et al. point out, this definition of security severely restricts the presence of non-determinism: for non-deterministic systems, a more intricate definition of non-interference will be necessary. Since FF is deterministic, we limit our attention in this paper to deterministic reactive systems, and the definition above satisfies our needs.

Definition 8.5. *A reactive system is deterministic if*

- *for all* $P \in ProducerState$ *the following holds:*

$$(P \xrightarrow{o} Q \wedge P \xrightarrow{o'} Q') \Rightarrow (o = o' \wedge Q = Q') \tag{2}$$

- *for all* $C \in ConsumerState$ *the following holds:*

$$(C \xrightarrow{i} P \wedge C \xrightarrow{i} P') \Rightarrow P = P' \tag{3}$$

The definitions in this section allow us to state our first goal for this paper more precisely: we want to construct an enforcement mechanism that ensures that FF is reactive non-interferent. Our non-interference proof will build on a result from Bohannon et al. [3]. They propose an interesting proof technique for establishing that a reactive system is non-interferent, based on the notion of *ID-bisimulation*.

Definition 8.6. *An ID-bisimulation on a reactive system is a label-indexed family of binary relations on states (written* $\sim_l$*) with the following properties:*

(a) if $Q \sim_l Q'$*, then* $Q' \sim_l Q$*;*

(b) if $C \sim_l C'$ *and* $C \xrightarrow{i} P$ *and* $C' \xrightarrow{i} P'$*, then* $P \sim_l P'$*;*

(c) if $C \sim_l C'$ *and* $\neg visible_l(i)$ *and* $C \xrightarrow{i} P$*, then* $P \sim_l C'$*;*

(d) if $P \sim_l C$ *and* $P \xrightarrow{o} Q$*, then* $\neg visible_l(o)$ *and* $Q \sim_l C$*;*

(e) if $P \sim_l P'$ *then either*

- $P \xrightarrow{o} Q$ *and* $P' \xrightarrow{o'} Q'$ *implies* $o = o'$ *and* $Q \sim_l Q'$*, or else*
- $P \xrightarrow{o} Q$ *implies* $\neg visible_l(o)$ *and* $Q \sim_l P'$*, or else*
- $P' \xrightarrow{o'} Q'$ *implies* $\neg visible_l(o')$ *and* $P \sim_l Q'$*.*

They show that the existence of an ID-bisimulation entails non-interference.

Theorem 8.1 ([3])**.** *If* $Q \sim_l Q$ *for all* l*, then* Q *is ID-secure.*

This theorem will be a key building block of our security proof: we will establish non-interference for our enforcement mechanism by showing the existence of an ID-bisimulation.

8.5 Proofs of the theorems

To prove the Theorem 4.1 first we need to prove the following lemma.

Lemma 8.1. *The wrapper over a deterministic reactive system is a deterministic reactive system.*

Proof. First we show that for every producer state (R, L) of a wrapper over a deterministic reactive system the following holds:

$$((R, L) \xrightarrow{o_1} (R_1, L_1) \wedge (R, L) \xrightarrow{o_2} (R_2, L_2)) \Rightarrow \\ (o_1 = o_2 \wedge (R_1, L_1) = (R_2, L_2))$$

$L = l :: L'$ is a list of levels. If $R(l) \xrightarrow{o} Q$ and $R(l) \xrightarrow{o'} Q'$ then $o = o'$ and $Q = Q'$ since the original system is deterministic. Hence, if $\mathrm{lbl}(o) = l$, then the rule [OUT-P] or [OUT-C] applies, otherwise [DROP-P] or [DROP-C] applies depending on whether Q is a consumer or a producer state. It means that $o_1 = o_2$ and $(R_1, L_1) = (R_2, L_2)$.

Now we show that for every consumer state $(R, \emptyset)$ of a wrapper over a deterministic reactive system the following holds:

$$((R, \emptyset) \xrightarrow{i} (R_1, L_1) \wedge (R, \emptyset) \xrightarrow{i} (R_2, L_2)) \Rightarrow \\ (R_1, L_1) = (R_2, L_2)$$

Since $(R, \emptyset)$ is a consumer state then only rule [LOAD] applies. Therefore, $(R_1, L_1) = (R_2, L_2)$. □

Lemma 4.1. *This l-similarity relation is an ID-bisimulation.*

Proof. Let us take two states of the wrapper (R_1, L_1) and (R_2, L_2). In the proof we will write first an original statement from the Definition 4.1 of ID-bisimulation and below we write the same statement in our notation.

a) if $Q_1 \sim_l Q_2$, then $Q_2 \sim_l Q_1$

if $(R_1, L_1) \approx_l (R_2, L_2)$ then $(R_2, L_2) \approx_l (R_1, L_1)$ obviously from the definition of l-similarity.

b) if $C_1 \sim_l C_2$ and $C_1 \xrightarrow{i} P_1$ and $C_2 \xrightarrow{i} P_2$, then $P_1 \sim_l P_2$

if $(R_1^C, \emptyset) \approx_l (R_2^C, \emptyset)$ and

$$(R_1^C, \emptyset) \xrightarrow{i} (R_1^P, L_1^P)$$
$$(R_2^C, \emptyset) \xrightarrow{i} (R_2^P, L_2^P)$$

then, $(R_1^P, L_1^P) \approx_l (R_2^P, L_2^P)$.

Only rule [LOAD] applies to the transitions because L component of the consumer states is empty. Since $R_1^C \approx_l R_2^C$ then for all $l' \leq l$: $R_1^C(l') = R_2^C(l')$. Then, $(R_1^P, L_1^P) \approx_l (R_2^P, L_2^P)$ because:

- Since the original reactive system is deterministic, for all $l' \leq l$, $R_1^C(l') \xrightarrow{i} P$ and $R_1^C(l') \xrightarrow{i} P'$ implies $P = P'$, therefore $R_1^P \approx_l R_2^P$,

- $L_1^P = L_2^P = Upper(i)$.

c) if $C_1 \sim_l C_2$ *and* $\neg visible_l(i)$ *and* $C_1 \xrightarrow{i} P_1$*, then* $P_1 \sim_l C_2$ if $(R_1^C, \emptyset) \approx_l (R_2^C, \emptyset)$ and $\text{lbl}(i) \not\leq l$ and

$$(R_1^C, \emptyset) \xrightarrow{i} (R_1^P, L_1^P)$$

then $(R_1^P, L_1^P) \approx_l (R_2^C, \emptyset)$.

Only rule [LOAD] applies to the transitions because L component of the consumer states is empty. Then, $(R_1^P, L_1^P) \approx_l (R_2^C, \emptyset)$ because:

- Since $R_1^C \approx_l R_2^C$ then for all $l' \leq l$: $R_1^C(l') = R_2^C(l')$. Since $\text{lbl}(i) \not\leq l$ then for all $l' \leq l$: $R_1^P(l') = R_1^C(l') = R_2^C(l')$, hence $R_1^C \approx_l R_2^C$.
- Since $\text{lbl}(i) \not\leq l$ and $L_1^P = Upper(i)$ only contains levels l' such that $\text{lbl}(i) \leq l'$ and $\text{lbl}(i) \not\leq l$, then $L_1^P|_l = \emptyset$.

d) if $P_1 \sim_l C_1$ *and* $P_1 \xrightarrow{o} Q_1$*, then* $\neg visible_l(o)$ *and* $Q_1 \sim_l C_1$
if $(R_1^P, L_1^P) \approx_l (R_1^C, \emptyset)$ and

$$(R_1^P, L_1^P) \xrightarrow{o} (R_1^Q, L_1^Q)$$

then ($\text{lbl}(o) \not\leq l$ or $o = \cdot$) and $(R_1^Q, L_1^Q) \approx_l (R_1^C, \emptyset)$ Since $(R_1^P, L_1^P) \approx_l (R_1^C, \emptyset)$, then $L_1^P|_l = \emptyset$. It means that L_1^P contains only levels l' such that $l' \not\leq l$.

If rule [DROP-P] or [DROP-C] is used, then $o = \cdot$. Otherwise the rule [OUT-P] or [OUT-C] is used, where $\text{lbl}(o)$ is in L_1^P, it means that $\text{lbl}(o) \not\leq l$. Hence, ($\text{lbl}(o) \not\leq l$ or $o = \cdot$).

Then, $(R_1^Q, L_1^Q) \approx_l (R_1^C, \emptyset)$ because

- Since $R_1^P \approx_l R_1^C$, then for all $l' \leq l$, $R_1^P(l') = R_1^C(l') \in ConsumerState$, so only for $l^* > l$ the state $R_1^P(l^*)$ is a producer state. Therefore, for all $l' \leq l$, $R_1^Q(l') = R_1^P(l') = R_1^C(l')$, so $R_1^Q \approx_l R_1^C$,
- According to all the rules from [OUT-P] to [DROP-C], L_1^Q is a suffix of L_1^P, therefore since $L_1^P|_l = \emptyset$, then $L_1^Q|_l = \emptyset$.

e) if $P_1 \sim_l P_2$ *then either*

1. $P_1 \xrightarrow{o_1} Q_1$ *and* $P_2 \xrightarrow{o_2} Q_2$ *implies* $o_1 = o_2$ *and* $Q_1 \sim_l Q_2$*, or else*
2. $P_1 \xrightarrow{o_1} Q_1$ *implies* $\neg visible_l(o_1)$ *and* $Q_1 \sim_l P_2$*, or else*
3. $P_2 \xrightarrow{o_2} Q_2$ *implies* $\neg visible_l(o_2)$ *and* $P_1 \sim_l Q_2$

if $(R_1^P, L_1^P) \approx_l (R_2^P, L_2^P)$ then either

1. $(R_1^P, L_1^P) \xrightarrow{o_1} (R_1^Q, L_1^Q)$ and $(R_2^P, L_2^P) \xrightarrow{o_2} (R_2^Q, L_2^Q)$ implies $o_1 = o_2$ and $(R_1^Q, L_1^Q) \approx_l (R_2^Q, L_2^Q)$, or else

2. $(R_1^P, L_1^P) \xrightarrow{o_1} (R_1^Q, L_1^Q)$ implies ($\text{lbl}(o_1) \not\leq l$ or $o_1 = \cdot$) and $(R_1^Q, L_1^Q) \approx_l (R_2^P, L_2^P)$, or else

3. $(R_2^P, L_2^P) \xrightarrow{o_2} R_2^Q, L_2^Q)$ implies ($\text{lbl}(o_2) \not\leq l$ or $o_2 = \cdot$) and $(R_1^P, L_1^P) \approx_l R_2^Q, L_2^Q)$

Since $(R_1^P, L_1^P) \approx_l (R_2^P, L_2^P)$ then $R_1^P \approx_l R_2^P$ and $L_1^P|_l = L_2^P|_l$.

Notice that $L_1^P \neq \emptyset$ and $L_2^P \neq \emptyset$ because (R_1^P, L_1^P) and (R_2^P, L_2^P) are producer states.

CASE e).1. $L_1^P|^l = \emptyset$ *and* $L_2^P|^l = \emptyset$
In this case $L_1^P|_l = L_2^P|_l = L_1^P = L_2^P = l_1, \ldots, l_m \neq \emptyset$ and $l_i \leq l$.
Then, one of the rules [OUT-P], [OUT-C], [DROP-P] or [DROP-C] applies and we prove condition e).1.
$R_1^P(l_1)$ and $R_2^P(l_1)$ are producer states because l_1 is in L_1^P and L_2^P.
If $\text{lbl}(o) = l_1$, then $R_1^P(l_1) \xrightarrow{o} Q$, $R_2^P(l_1) \xrightarrow{o'} Q'$. Since the original reactive system is deterministic, then $o = o'$ and $Q = Q'$.
If $Q \in \mathit{ProducerState}$ then [OUT-P] applies to both (R_1^P, L_1^P) and (R_2^P, L_2^P). Hence

$$(R_1^P, L_1^P) \xrightarrow{o_1} (R_1^Q, L_1^Q)$$
$$(R_2^P, L_2^P) \xrightarrow{o_2} (R_2^Q, L_2^Q)$$

obviously implies $o_1 = o_2 = o$ and $(R_1^Q, L_1^Q) \approx_l (R_2^Q, L_2^Q)$.
In the same way we prove the cases of rules [OUT-C], [DROP-P] and [DROP-C].

To ease the proof of cases e).2 and e).3 we will use parameters i, j and substitution $i = 1, j = 2$ proves e).2 while substitution $i = 2, j = 1$ proves e).3.
CASE e).2, e).3. $L_i^P|^l \neq \emptyset$
Notice that we don't specify the condition on the $L_j^P|^l$, for this case it does not matter.
According to definition, $L_i^P|^l = l'_1, \ldots, l'_k$ and $l'_j \not\leq l$. Then for the transition

$$(R_i^P, L_i^P) \xrightarrow{o_i} (R_i^Q, L_i^Q)$$

one of the rules from [OUT-P] to [DROP-C] applies. If the rule [DROP-P] or [DROP-C] applies, then $o_i = \cdot$. If the rule [OUT-P] or [OUT-C] applies, then the first element l'_1 of L_i^P that is taken is such that $l'_1 \not\leq l$ and hence $\text{lbl}(o_i) = l'_1 \not\leq l$.
Then, $(R_i^Q, L_i^Q) \approx_l (R_j^P, L_j^P)$ because

- According to all the rules [OUT-P], [OUT-C], [DROP-P] or [DROP-C], R_i^P is changed only at level $l'_1 \not\leq l$, then $R_i^Q \approx_l R_i^P \approx_l R_j^P$,
- $L_i^Q|_l = L_i^P|_l = L_j^P|_l$ because in [OUT-C] and [DROP-C] only level l'_1 such that $l'_1 \not\leq l$ is excluded from L_i^P to get L_i^Q, and in [OUT-P] or [DROP-P], $L_i^Q = L_i^P$.

□

Theorem 4.1 (Security)**.** *All the states of the wrapper are ID-secure.*

Proof. Since for every state (R, L) of the wrapper, $(R, L) \approx_l (R, L)$ and $\approx_l$ is an ID-bisimulation according to Lemma 4.1, then (R, L) is secure according to the [3, Theorem 4.5]. □

Lemma 4.2. *If* $O \approx_l^{ID} O'$*, then* observer-indistinguishable$_{l'}(O, O')$ *for all* $l' \leq l$*.*

Proof. Because of coinduction, it suffices to prove for any $l' \leq l$ that if $O \approx_l^{ID} O'$ then either

- $O = o :: O''$, $lbl(o) \neq l'$ and $O'' \approx_l^{ID} O'$
- $O' = o' :: O''$, $lbl(o') \neq l'$ and $O \approx_l^{ID} O''$
- $O = o :: O''$, $O' = o :: O'''$, $lbl(o') = l'$ and $O'' \approx_l^{ID} O'''$
- $O = O' = []$.

We can prove this by case analysis. Either both lists are empty, or (by symmetry), $O = o :: O''$. Then either $lbl(o) \leq l$ or not. If not, then automatically $lbl(o) \neq l'$ and we can verify that $O'' \approx_l^{ID} O$. In the case that $lbl(o) \leq l$, this means that $O \rhd_l o :: O''$. Because $O \approx_l^{ID} O'$, O' cannot be the empty list, so $O' = o' :: O'''$. If $lbl(o') \leq l$, then $o = o'$, and $O'' \approx_l^{ID} O'''$. If $lbl(o') \not\leq l$, then automatically $lbl(o') \neq l'$ and we can verify that $O \approx_l^{ID} O'''$. □

To prove the Theorem 4.2 and Lemma 4.3 we need to give several additional definitions and lemas.

Definition 8.7. *For a finite input stream* I*, a reactive system* Q*, finite output stream* O *and state* C*, we define* $Q(I) \stackrel{O}{\Rightarrow}_t C$ *with the following rules:*

$$C([]) \stackrel{[]}{\Rightarrow}_t C \qquad \frac{C \stackrel{i}{\to} P \qquad P(I) \stackrel{O}{\Rightarrow}_t C'}{C(i :: I) \stackrel{O}{\Rightarrow}_t C'} \qquad \frac{P \stackrel{o}{\to} Q \qquad Q(I) \stackrel{O}{\Rightarrow}_t C'}{P(I) \stackrel{o::O}{\Rightarrow}_t C'}$$

Lemma 8.2. *For* I_0 *and* O *finite, we have that* $Q(I_0) = O$ *iff there exists a* C' *such that* $Q(I_0) \stackrel{O}{\Rightarrow}_t C'$*.*

Lemma 8.3. *For a reactive system* Q*, an input stream* I *such that* $I = \text{concat}(I_0, I')$ *with* I_0 *finite, an output stream* O *such that* $O = \text{concat}(O_0, O')$ *with* O_0 *finite, we have that* $Q(I) = O$ *iff there exists a* Q' *such that* $Q(I_0) \stackrel{O_0}{\Rightarrow}_t Q'$ *and* $Q'(I') = O'$*.*

Lemma 8.4. *For all* I_0*, we have that* $I_0 \approx_l^{ID} \pi_l(I_0)$*.*

Lemma 8.5. *Suppose* $W = (R, L)$ *with* $R(l) = Q$ *and* $l \in L$ *iff* $Q \in ProducerState$. *Suppose* I_0 *is finite and* $W(I_0) \stackrel{O_0}{\Rightarrow}_t W'$ *with* $W' = (R', L')$, $R'(l) = S'$. *Then* $Q(\pi_l(I_0)) \stackrel{O'_0}{\Rightarrow}_t S'$ *with* observer-indistinguishable$_l(O_0, O'_0)$ *and* $S' \in ProducerState$ *iff* $l \in L'$.

Proof. We prove this by induction on the (finite) derivation of $W(I_0) \stackrel{O}{\Rightarrow}_t W'$.

Suppose W is in a consumer state. Then so must be Q. Suppose $I_0 = [] = \pi_l(I_0)$ and $L = \phi$. Then $O = Q(I_0) = [] = W(I_0) = O'$, $W' = W$ and $R' = R$.

Suppose $I_0 = i :: I'_0$ and $L = \phi$. If $\text{lbl}(i) \leq l$, then $\pi_l(I_0) = i :: \pi_l(I'_0)$. We know that $W \stackrel{i}{\rightarrow} W''$ with $W'' = (R'', L'')$, $R''(l) = P$, $C \stackrel{i}{\rightarrow} P$, $l \in L''$ and $W''(I'_0) \stackrel{O}{\Rightarrow}_t W'$. By induction, we then know that $P(\pi_l(I'_0)) \stackrel{O'}{\Rightarrow}_t S'$ and $S' \in ProducerState$ iff $l \in L'$ with observer-indistinguishable$_l(O, O')$ and thus $Q(\pi_l(I_0)) \stackrel{O'}{\Rightarrow}_t S'$.

The other cases are similar. □

Lemma 8.6. *Assume a given security level* l *and input stream* I. *Suppose reactive system state* Q_0 *behaves securely for input* I *and wrapper state* $W_0 = (R_0, L_0)$ *is such that* $R_0(l) = Q_0$, $l \in L_0$ *iff* $Q_0 \in ProducerState$. *Suppose* $Q_0(I_0) \stackrel{O_0}{\Rightarrow}_t Q$, $W_0(I_0) \stackrel{O'_0}{\Rightarrow}_t W = (R, L)$ *and* observer-indistinguishable$_l(O_0, O'_0)$. *Suppose* $Q(I') = O_Q$ *and* $W(I') = O_W$ *with* $I = \text{concat}(I_0, I')$. *Then we have* observer-indistinguishable$_l(O_Q, O_W)$.

Proof. Because of coinduction, it suffices to prove that with the assumptions in the theorem statement, either $O_Q = O_W = []$ or there exist O'_Q, O'_W, Q', I'_0, I', O''_0, O'''_0, $W' = (R', L')$ such that $Q_0(I'_0) \stackrel{O''_0}{\Rightarrow}_t Q'$, $W_0(I'_0) \stackrel{O'''_0}{\Rightarrow}_t W'$, observer-indistinguishable$_l(O_0, O'_0)$ and $Q'(I') = O'_Q$, $W'(I') = O'_W$ and $I = \text{concat}(I_0, I')$ and one of the following holds:

1. $O_Q = o :: O'_Q$, $lbl(o) \neq l$, $O_W = O'_W$.
2. $O_W = o :: O'_W$, $lbl(o) \neq l$, $O_Q = O'_Q$.
3. $O_Q = o :: O'_Q$, $O_W = o :: O'_W$.

We can prove this by case analysis.

- If W is in a producer state then $L = l' :: L'$ and there exists an o such that $W \stackrel{o}{\rightarrow} W'$, $W' = (R', L')$. We distinguish two sub-cases:
  - $lbl(o) \neq l$. Then we know that $l' \neq l$ or $o = \cdot$. We can take $Q' = Q$, $I'_0 = I_0$, $I' = I$, $O''_0 = O_0$, $O'''_0 = \text{concat}(O'_0, [o])$, and case 2 above applies.
  - $lbl(o) = l$. Then we know that $W_0(I_0) \stackrel{\text{concat}(O'_0,[o])}{\Rightarrow}_t W'$. Therefore, lemma 8.5 above tells us that $Q(\pi_l(I_0)) \stackrel{O_{Q,\pi}}{\Rightarrow}_t Q_\pi$ and states that observer-indistinguishable$_l(\text{concat}(O'_0, [o]), O_{Q,\pi})$. Because of

lemma 8.3, $O_{Q,\pi}$ must be a prefix of $Q_0(\pi_l(I_0))$ and we know that observer-indistinguishable$_l(Q_0(\pi_l(I_0)), Q_0(I_0))$. This means that Q cannot be in a consumer state, because then $Q_0(I_0) = O_0$, which cannot be observer-indistinguishable$_l$ from $Q_0(\pi_l(I_0))$. Therefore $Q \xrightarrow{o'} Q'$. We now distinguish two possibilities:

  * $lbl(o') \neq l$. Then case 1 above applies with $I'_0 = I_0$, $I' = I$, $O''_0 = o :: O_0$, $O'''_0 = O'_0$.
  * $lbl(o') = l$. We know (because of lemma 8.3), that concat$(O_0, [o'])$ is a prefix of $Q_0(I_0)$ and since the following predicate holds observer-indistinguishable$_l(Q_0(\pi_l(I_0)), Q_0(I_0))$, $o = o'$. Therefore, case 3 above applies with $I'_0 = I_0$, $I' = I$, $O''_0 = o :: O_0$, $O'''_0 = o :: O'_0$.

- If W is in a consumer state, then $W_0(I_0) = O'_0$. We distinguish two cases:
  - Q is a producer state. In this case, there exists an o such that $Q \xrightarrow{o} Q'$. Therefore, $Q_0(I_0) \overset{\text{concat}(O_0,[o])}{\Rightarrow_t} Q'$. Then we know (because of lemma 8.3) that concat$(O_0, [o])$ is a prefix of $Q_0(I_0)$. On the other hand, because $W_0(I_0) = O'_0$, lemma 8.5 tells us that $Q_0(\pi_l(I_0)) \overset{O_{Q,\pi}}{\Rightarrow_t} Q_\pi$ with observer-indistinguishable$_l(O_{Q,\pi}, O'_0)$ and Q_π is a consumer state. Therefore, $Q_0(\pi_l(I_0)) = O_{Q,\pi}$. Furthermore, we know that observer-indistinguishable$_l(Q_0(I_0), Q_0(\pi_l(I_0)))$. Because concat$(O_0, [o])$ is a prefix of $Q_0(I_0)$ yet observer-indistinguishable$_l(O_{Q,\pi}, O'_0)$ and observe-indistinguishable$_l(O_0, O'_0)$ hold, it must be the case that $lbl(o) \neq l$. Thus, case 1 applies with $I'_0 = I_0$, $I' = I$, $O''_0 = \text{concat}(O_0, [o])$, $O'''_0 = O'_0$ and $W' = W$.
  - Q is a consumer state. Then either $I' = []$, $O_Q = Q(I') = []$ and $O_W = W(I') = []$. In the other case, we can replace I_0 with concat$(I_0, [\text{head}\, I'])$, I' with tail I' and Q and W with Q_p and W_p such that $Q \overset{\text{head}\, I'}{\rightarrow} Q_p$ and $W \overset{\text{head}\, I'}{\rightarrow} W_p$ and apply the case above that W is in a producer state.

□

Theorem 4.2 (Precision for individual runs)**.** *Suppose a given reactive system state Q behaves securely for input I and $Q(I) = O_Q$. Define the corresponding wrapper $W = (R_Q, L)$ with $R_Q(l) = Q$ for all l, $L = \emptyset$ if $Q \in ConsumerState$ and $L = \mathcal{L}$ if $Q \in ProducerState$. For $O_W = W(I)$, we have that $O_Q \approx^{obs} O_W$.*

Proof. This is easily proven by applying lemma 8.6 for an arbitrary l with $I_0 = O_0 = O'_0 = []$. □

Lemma 4.3. *If a given reactive system state Q is ID-secure, then it behaves securely for any input I.*

Proof. This follows from the definitions using lemmas 4.2 and 8.4. □

Theorem 4.3 (Precision)**.** *Suppose a given reactive system state Q is ID-secure, and $Q(I) = O$. Define the corresponding wrapper $W = (R_Q, L)$ with $R_Q(l) = Q$ for all l, $L = \emptyset$ if $Q \in ConsumerState$ and $L = \mathcal{L}$ if $Q \in ProducerState$. For $O' = W(I)$, we have that $O \approx^{obs} O'$.*

Proof. This follows directly from Theorem 4.2 and Lemma 4.3. □

Lemma 6.1. *Suppose the current state of the wrapped system is $W = (R, \emptyset)$. If R is defined for l and l' and the input I is such that $\pi_l(I) = \pi_{l'}(I)$ (the input I looks the same at l and l'), then $R(l) = R(l')$.*

Proof. We prove the lemma by induction on the length of I.

- Initial state of the wrapped system is $(R, \emptyset)$ such that for all levels l, $R(l)$ points to an initial state of the original browser. Hence initially $R(l) = R(l')$.
- Let us assume that on the input stream I for some state $W = (R, L)$ of the wrapped system $R(l) = R(l')$.
- Let us prove that on the input stream I after one transition of the wrapped system
$$(R, L) \xrightarrow{s} (R', L') \tag{4}$$
we have $R'(l) = R'(l')$.

  In case transition (4) corresponds to the rule [LOAD] and the new input event i arrived, then either $\text{lbl}(i) \leq l$ or not. In case it's not, $\pi_l(I)$ does not contain i and so $\pi_{l'}(I)$ does not contain it either. Hence, $\text{lbl}(i) \not\leq l'$, and it means that $R'(l) = R'(l') = R(l) = R(l')$. In case $\text{lbl}(i) \leq l$ by similar reasoning we have that $\text{lbl}(i) \leq l'$, therefore according to the [LOAD] rule, $R'' = R[l \mapsto P]$ and $R' = R''[l' \mapsto P]$, therefore $R'(l) = R'(l') = P$. In case transition (4) corresponds to any other of the rules, $R(l) \xrightarrow{o} P$ and $R(l') \xrightarrow{o'} P'$. Since the original browser is deterministic, $P = P'$ and so $R'(l) = R'(l') = P$.

□