

Modular Monitor Extensions for Information Flow Security in JavaScript

José Fragoso Santos, Ana Almeida Matos and Tamara Rezk

Abstract

Client-side JavaScript programs often interact with the web page into which they are included, as well as with the browser itself, through APIs such as the DOM API, the XMLHttpRequest API, and the W3C Geolocation API. Precise reasoning about JavaScript security must therefore take API invocation into account. However, the continuous emergence of new APIs, and the heterogeneity of their forms and features, renders API behavior a moving target that is particularly hard to capture. To tackle this problem, we propose a methodology for modularly extending sound JavaScript information flow monitors with a generic API. Hence, to verify whether an extended monitor complies with the proposed noninterference property requires only to prove that the API satisfies a predefined set of conditions. In order to illustrate the practicality of our methodology, we show how an information flow monitor-inlining compiler can take into account the invocation of arbitrary APIs, without changing the code or the proofs of the original compiler. We provide an implementation of such a compiler with an extension for handling a fragment of the DOM Core Level 1 API. Furthermore, our implementation supports the addition of monitor extensions for new APIs at runtime.

1 Introduction

Isolation properties guarantee protection for trusted JavaScript code from malicious code. The noninterference property [11] provides the mathematical foundations for reasoning precisely about isolation. In particular, noninterference properties guarantee absence of flows from confidential/untrusted resources to public/trusted ones.

Although JavaScript can be used as a general-purpose programming language, many JavaScript programs are designed to be executed in a browser in the context of a web page. Such programs often interact with the web page in which they are included, as well as with the browser itself, through Application Programming Interfaces (APIs). Some APIs are fully implemented in JavaScript, whereas others are built with a mix of different technologies, which can be exploited to conceal sophisticated security violations. Thus, understanding the behavior of client-side web applications, as well as proving their compliance with a given security policy, requires cross-language reasoning. The size, complexity, and number of commonly used APIs poses an important challenge to any attempt at formally reasoning about the security of JavaScript programs [15]. To tackle this problem, we propose a methodology for extending JavaScript monitored semantics. This methodology allows us to verify whether a monitor complies with the proposed noninterference property in a modular way. Thus, we make it possible to prove that a security monitor is still noninterferent when extending it with a new API, without having to revisit the whole model. Generally, an API can be viewed as a particular set of specifications that a program can follow to make use of the resources provided by another particular application. For client-side JavaScript programs, this definition of API applies both to: (1) interfaces of services that are provided to the program by the environment in which it executes, namely the web browser (for instance, the DOM, the XMLHttpRequest, and the W3C Geolocation APIs); (2) interfaces of JavaScript libraries that are explicitly included by the programmer (for instance, jQuery, Prototype.js, and Google Maps Image API). In the context of this work, the main difference between these two types of APIs is that in the former case their semantics escapes the JavaScript semantics, whereas in the latter it does not. The methodology proposed here was designed



licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics



LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

as a generic way of extending security monitors to deal with the first type of APIs. Nevertheless, we can also apply it to the second type whenever we want to execute the library's code in the original JavaScript semantics instead of the monitored semantics.

► **Example 1** (Running example: A Queue API). Consider the following API for creating and manipulating priority queues. The API is available to the programmer through the global variable *queueAPI*, and variable *queueObj* is bound to a concrete queue:

```
queueAPI.queue(): creates a new priority queue;
queueObj.push(el, priority): adds a new element to the queue;
queueObj.pop(): pops the element with the highest priority.
```

The method calls from this API cannot be verified by the JavaScript monitor, as we are assuming that the code of the methods is not available to the JavaScript engine. Furthermore, the specification of the queue API may not obey the JavaScript semantics and hence prevention of the security leaks may need different constraints.

In order to extend a JavaScript security monitor to control the behavior of this API, one has to define what we call an *API Register* to set the security constraints associated to the corresponding API method calls on *queueAPI* and *queueObj*. API method calls should be implemented as interception points of the monitor semantics and the API Register should then make the invocation of these methods if the security constraints are satisfied.

The following questions then arise: What constraints must we impose on the new API register in order to preserve the noninterference guarantees of the JavaScript monitor? Is it possible to modularly prove noninterference of the extended monitor without revisiting the whole set of constraints, including those of the JavaScript monitor?

There are two main approaches for implementing a monitored JavaScript semantics: either one modifies a JavaScript engine so that it also implements the security monitor (as in [16]), or one inlines the monitor in the original program (as in [17], [10], and [24]). We show how to extend an information flow monitor-inlining compiler so that it also takes into account the invocation of APIs. Our extensible compiler requires each API to be associated with a set of JavaScript methods that we call its *IFlow Signature*, which describes how to handle the information flows triggered by its invocation. We provide a prototype of the compiler, which is available online [23]. A user can easily extend it by loading new IFlow signatures. Using the compiler, we give realistic examples of how to prevent common security violations that arise from the interaction between JavaScript and the DOM API.

The contributions of the paper are: (1) a methodology for extending JavaScript monitors with API monitoring (Section 3.2), (2) the design of an extensible information flow monitor-inlining compiler that follows our methodology (Section 4), (3) an implementation [23] of a JavaScript information flow monitor-inlining compiler (Section 5) that handles an important subset of the DOM API and is extensible with new APIs by means of IFlow Signatures.

2 Related Work

We refer the reader to a recent survey [6] on web scripts security and to [22] for a complete survey on information flow enforcement mechanisms up to 2003, while focusing here on the most closely related work on dynamic mechanisms for enforcing noninterference. Flow-sensitive monitors for enforcing noninterference can be broadly divided into two classes: those that are purely dynamic, such as [2], [3], and [4], and those that mix runtime monitoring with static analysis, commonly referred to as *hybrid monitors*, such as [27], [14], and [25]. In contrast to hybrid monitors, which rely on static analysis to reason about implicit flows

that arise due to untaken execution paths, purely dynamic monitors do not rely on any kind of static analysis. The authors of [2], [3], and [4] propose three alternative strategies in designing sound purely dynamic information flow monitors. The *no-sensitive-upgrade* strategy forbids the update of public resources inside private contexts. The *permissive-upgrade* strategy allows sensitive upgrades to take place, but marks the resources upgraded in sensitive contexts and forbids the program to branch depending on the content of these resources. Finally, the *multiple facet* strategy surpasses the limitations of the first two (which can potentially abort the execution of secure programs) by the use of multiple faceted values. The intuition behind this strategy is that values must appear differently to observers at different security levels: the security monitor simulates multiple executions for different security levels. Russo et. al [20] show that purely dynamic monitors always reject executions that would have been accepted using static enforcement mechanisms. In this work, we show how to extend an arbitrary information flow monitor without making any assumption on its particular design. However, the monitor extensions that illustrate the applicability of our methodology follow the no-sensitive-upgrade discipline of [2].

Chudnov and Naumann [10] propose the inlining of a hybrid information flow monitor for a simple imperative language. Simultaneously, Magazinius et al. [17] also propose a monitor inlining transformation for an imperative language, with the novel feature of performing inlining on the fly to allow for dynamic code evaluation.

Hedin and Sabelfeld [16] are the first to propose a runtime monitor for enforcing non-interference for JavaScript. The technique that we present for extending security monitors can be applied to this monitor, which is a purely dynamic monitor and therefore liable to the limitations described in [20]. To account for these limitations, Birgisson et. al [8] show how to use tests in order to boost the permissiveness of [16]. Each time a security error arises during a test, the program is modified with an annotation that prevents the same error from reoccurring. In [24], we present the inlining of an information flow monitor for a realistic subset of JavaScript. In the implementation presented in this paper we extend the inlined monitor of [24] with the DOM API, applying the methodology proposed here. Despite targeting JavaScript, these two works, as our own, do not model the reactive aspect of client-side web applications. Bohannon et al. [9] present a definition of noninterference for reactive programs such as web scripts as well as a runtime monitor for enforcing it. Devriese et al. [12] propose a dynamic enforcement mechanism for reactive noninterference called secure multi execution and implement it in the Featherweight Firefox browser [7].

Taly et al. [26] study API confinement. They provide a static analysis designed to verify whether, when integrating the code of an API in an arbitrary page, the integrator code cannot interact with the API and cause it to leak its confidential resources. They use a more restrictive notion of a web API than the one used in this work, in the sense that they use the term API only to refer to JavaScript libraries whose code is explicitly included by the programmer and, therefore, available for either runtime or static analysis.

Russo et al. [21] present an information flow monitor for a WHILE language with primitives for manipulating DOM-like trees and prove it sound. They do not model references. Frago Santos et al. [18] model references and show that they entail new information leaks, in a simple language that models the DOM API, based on the work of Gardner et al. [13]. We apply techniques of [18] to our implementation of the DOM API extension based on the specification of DOM Core Level 1 API [19]. Recent work [28] presents an information flow monitor for JavaScript extended with the DOM API that also considers event handling loops. To the best of our knowledge, no prior work proposes a generic methodology to extend JavaScript monitors and inlined compilers with arbitrary web APIs.

3 Modular Extensions for JavaScript Monitors

In this section we show how to extend a noninterferent monitor so that it takes into account the invocation of web APIs, while preserving the noninterference property.

3.1 Noninterferent JavaScript Monitors

JavaScript Memory Model. In JavaScript [1], objects can be seen as partial functions mapping strings to values. The strings in the domain of an object are called its properties. Memories are mappings from references to objects. In the following, we assume that memories include a reference to a special object called the *global object* pointed to by a fixed reference $\#glob$, that binds global variables. In this presentation, objects, properties, memories, references and values, are ranged over by o, p, μ, r and v , respectively.

We use the notation $[p_0 \mapsto v_0, \dots, p_n \mapsto v_n]$ for the partial function that maps p_i to v_i where $i = 0, \dots, n$, and $f[p_0 \mapsto v_0, \dots, p_n \mapsto v_n]$ for the function that coincides with f everywhere except in p_0, \dots, p_n , which are otherwise mapped to v_0, \dots, v_n respectively. Furthermore, we denote by $\text{dom}(f)$ the domain of a function f , and by $f|_P$ the restriction of the domain of f to P (when $P \subseteq \text{dom}(f)$). Finally, we write $f(r)(p)$ instead of $(f(r))(p)$, the application of the image of r by function f to p .

Sequences are denoted by stacking an arrow as in \vec{v} , and ϵ denotes the empty sequence. The length of \vec{v} is given by $|\vec{v}|$ and \cdot denotes concatenation of sequences.

Security Setting. Information flow policies such as noninterference are specified over security labelings that assign security levels, taken from a given security lattice, to the observable resources of a program. In the following, we use a fixed lattice \mathcal{L} of security levels ranged over by σ . We denote by \leq its order relation, by $\sigma_0 \sqcup \sigma_1$ the least upper bound between levels σ_0 and σ_1 , and by $\sqcup \vec{\sigma}$ the least upper bound of all levels in the sequence $\vec{\sigma}$. In the examples, we consider two security levels $\{H, L\}$ such that $L \leq H$, meaning that resources labeled with *high* level H are at least as confidential as those labeled with *low* level L .

In our setting, a security labeling is as a pair $\langle \Gamma, \Sigma \rangle$, where Γ maps references, followed by properties, to security levels, and Σ maps references to security levels. Then, given an object o pointed to by a reference r , if defined, $\Gamma(r)(p)$ corresponds to the security levels associated with o 's property p , and $\Sigma(r)$ with o 's domain. The latter, also referred to as o 's *structure security level*, controls the observability of the existence of properties [16].

We say that memory μ is well-labeled by $\langle \Gamma, \Sigma \rangle$ if $\text{dom}(\Gamma) = \text{dom}(\Sigma) \subseteq \text{dom}(\mu)$ and for every reference $r \in \text{dom}(\Gamma)$, $\text{dom}(\Gamma(r)) \subseteq \text{dom}(\mu(r))$.

Security Monitor. JavaScript programs are statements, that include expressions, ranged over by s and e , respectively. We model an information flow monitor as a small-step semantics relation \rightarrow_{IF} between configurations of the form $\langle \mu, s, \vec{\sigma}_{\text{pc}}, \Gamma, \Sigma, \vec{\sigma} \rangle$ composed of (1) a memory μ (2) a statement s , that is to execute, (3) a sequence of security levels $\vec{\sigma}_{\text{pc}}$, matching the expressions on which the original program branched to reach the current context, (4) a security labeling $\langle \Gamma, \Sigma \rangle$, and (5) a sequence of security levels $\vec{\sigma}$ matching the reading effects of the subexpressions of the expression being computed.

The *reading effect* [22] of an expression is defined as the least upper bound on the security levels of the resources on which the value to which it evaluates depends. Additionally, we assume that the reading effect of an expression is always higher than or equal to the level of the context in which it is evaluated, $\sqcup \vec{\sigma}_{\text{pc}}$.

Low-equality. In order to account for a non-deterministic memory allocator, we rely on a partial injective function β which relates observable references that point to the same

resource in different executions of the same program [5]. The β relation is extended to relate observable values via the β -equality, which is denoted \sim_β : two objects are β -equal if they have the same domain and all their corresponding properties are β -equal; primitive values and parsed functions are β -equal if syntactically equal; and, two references r_0 and r_1 are β -equal if the latter is the image by β of the former.

Two memories μ_0 and μ_1 are said to be *low-equal* with respect to labelings $\langle \Gamma_0, \Sigma_0 \rangle$ and $\langle \Gamma_1, \Sigma_1 \rangle$, a security level σ , and a partial injective function β , written $\mu_0, \Gamma_0, \Sigma_0 \approx_{\beta, \sigma} \mu_1, \Gamma_1, \Sigma_1$, if μ_0 and μ_1 are well-labeled by $\langle \Gamma_0, \Sigma_0 \rangle$ and $\langle \Gamma_1, \Sigma_1 \rangle$ respectively, and for all references $r_0, r_1 \in \text{dom}(\beta)$, such that $r_1 = \beta(r_0)$, the following hold:

1. The observable domains (i.e. set of observable properties) of the objects pointed by r_0, r_1 , coincide: $P_\sigma = \{p \in \text{dom}(\mu_0(r_0)) \mid \Gamma_0(r_0)(p) \leq \sigma\} = \{p \in \text{dom}(\mu_1(r_1)) \mid \Gamma_1(r_1)(p) \leq \sigma\}$;
2. The objects pointed by r_0, r_1 coincide in their observable domain: $\mu_0(r_0)|_P \sim_\beta \mu_1(r_1)|_P$;
3. If the structure security level of either object pointed by r_0, r_1 is observable ($\Sigma_0(r_0) \leq \sigma$ or $\Sigma_1(r_1) \leq \sigma$), then their domains entirely coincide: $\text{dom}(\mu_0(r_0)) = \text{dom}(\mu_1(r_1))$.

We extend informally the definition of low-equality to sequences of labeled values and to program continuations (the interested reader can find the formal definitions in [23]). Two sequences of labeled values are low-equal with respect to a given security level σ , denoted $\vec{v}_0, \vec{\sigma}_0 \approx_{\beta, \sigma} \vec{v}_1, \vec{\sigma}_1$ if for each position of both sequences, either the two values in that position are low-equal, or the levels that are associated with both of them are not observable. Low-equality between program continuations $s_0, \vec{\sigma}_{pc0}, \vec{\sigma}_0 \approx_{\beta, \sigma} s_1, \vec{\sigma}_{pc1}, \vec{\sigma}_1$ relaxes syntactic equality between programs in order to relate the intermediate states of the execution of the same original program in two low-equal memories, as illustrated by the following example.

► **Example 2 (Low-equal program continuations).** Consider the program $x = y$, an initial labeling $\langle \Gamma, \Sigma \rangle$ such that $\Gamma(\#glob)(x) = \Gamma(\#glob)(y) = H$, and two memories μ_0 and μ_1 such that $\mu_i = [\#glob \mapsto [x \mapsto \text{undefined}, y \mapsto i]]$, for $i \in \{0, 1\}$. The execution of one computation step of this program in μ_0 and μ_1 yields the programs $x = 0$ and $x = 1$. Since the reading effects associated with the values 0 and 1 are both H , the expressions $x = 0$ and $x = 1$ are low-equal. Formally: $x = 0, \langle L \rangle, \langle H \rangle \approx_{\text{id}, L} x = 1, \langle L \rangle, \langle H \rangle$ (where id is the identity function).

Finally, two monitor configurations $\langle \mu_0, s_0, \vec{\sigma}_{pc0}, \Gamma_0, \Sigma_0, \vec{\sigma}_0 \rangle$ and $\langle \mu_1, s_1, \vec{\sigma}_{pc1}, \Gamma_1, \Sigma_1, \vec{\sigma}_1 \rangle$ are said to be *low-equal* w.r.t a level σ and function β , written $\langle \mu_0, s_0, \vec{\sigma}_{pc0}, \Gamma_0, \Sigma_0, \vec{\sigma}_0 \rangle \approx_{\beta, \sigma} \langle \mu_1, s_1, \vec{\sigma}_{pc1}, \Gamma_1, \Sigma_1, \vec{\sigma}_1 \rangle$, if $\mu_0, \Gamma_0, \Sigma_0 \approx_{\beta, \sigma} \mu_1, \Gamma_1, \Sigma_1$ and $s_0, \vec{\sigma}_{pc0}, \vec{\sigma}_0 \approx_{\beta, \sigma} s_1, \vec{\sigma}_{pc1}, \vec{\sigma}_1$.

Noninterferent Monitor. In the remaining of the paper, we consider only *noninterferent* JavaScript monitors. As usual, a monitor \rightarrow_{IF} is noninterferent, written $\mathbf{NI}_{\text{mon}}(\rightarrow_{\text{IF}})$, if its application on two low-equal configurations produces two low-equal configurations.

► **Definition 3 (Monitor Noninterference).** A monitor \rightarrow_{IF} is said to be *noninterferent*, written $\mathbf{NI}_{\text{mon}}(\rightarrow_{\text{IF}})$, if for every programs s_0, s_1 , memories μ_0, μ_1 , and labeling $\langle \Gamma, \Sigma \rangle$, such that μ_0, μ_1 are well-labeled by $\langle \Gamma, \Sigma \rangle$ and, for all security levels σ , there exists β such that $\langle \mu_0, s_0, \epsilon, \Gamma, \Sigma, \epsilon \rangle \approx_{\beta, \sigma} \langle \mu_1, s_1, \epsilon, \Gamma, \Sigma, \epsilon \rangle$, if $\langle \mu_0, s_0, \epsilon, \Gamma, \Sigma, \epsilon \rangle \rightarrow_{\text{IF}}^* \langle \mu'_0, v'_0, \epsilon, \Gamma', \Sigma', \sigma' \rangle$ and $\langle \mu_1, s_1, \epsilon, \Gamma, \Sigma, \epsilon \rangle \rightarrow_{\text{IF}}^* \langle \mu'_1, v'_1, \epsilon, \Gamma', \Sigma', \sigma' \rangle$ then there is an extension β' of β such that $\langle \mu'_0, v'_0, \epsilon, \Gamma', \Sigma', \sigma' \rangle \approx_{\beta, \sigma} \langle \mu'_1, v'_1, \epsilon, \Gamma', \Sigma', \sigma' \rangle$.

3.2 API Extensions to JavaScript Monitors

API relation. An API is modeled as a semantic relation $\Downarrow_{\text{API}}^{\text{JS}}$ of the form $\langle \mu, \vec{v} \rangle \Downarrow_{\text{API}}^{\text{JS}} \langle \mu', v' \rangle$ where μ is the memory in which the API is executed, μ' is the resulting memory, the sequence

of values \vec{v} corresponds to the arguments of the API invocation, and v' is the value to which the API invocation evaluates. Accordingly, a *monitored API relation*, \Downarrow_{API} , has the form

$$\langle \mu, \Gamma, \Sigma, \vec{v}, \vec{\sigma} \rangle \Downarrow_{\text{API}} \langle \mu', \Gamma', \Sigma', v, \sigma \rangle$$

which adds to the original API configuration the initial and final labelings $\langle \Gamma, \Sigma \rangle$ and $\langle \Gamma', \Sigma' \rangle$ (respectively), the sequence of security levels $\vec{\sigma}$ that is associated with the arguments of the API invocation, and their corresponding reading effect σ .

API register. The bridge between API invocations and the corresponding monitored API semantics is performed by a *API register*, denoted by \mathcal{R}_{API} . We define an API register as a function that, given a memory and a sequence of values, returns a monitored API relation.

► **Example 4 (Queue API Register).** In order for an extended monitor to take into account the methods of the Queue API from Example 1, the API Register must be extended to handle invocations of the Queue API methods. In the following, \Downarrow_{QU} , \Downarrow_{PU} , and \Downarrow_{PO} are the API relations corresponding to each one of the methods of the Queue API:

$$\mathcal{R}_Q(\mu, \langle r, m, \dots \rangle) = \begin{cases} \Downarrow_{QU} & \text{if } m = \text{“queue”} \wedge \$q \in \text{dom}(\mu(r)) \\ \Downarrow_{PU} & \text{if } m = \text{“push”} \wedge \$q \in \text{dom}(\mu(r)) \\ \Downarrow_{PO} & \text{if } m = \text{“pop”} \wedge \$q \in \text{dom}(\mu(r)) \end{cases}$$

The idea is to “mark” the Queue API object (the one bound to variable *queueAPI*) as well as the concrete queue objects, with a special property (in this case, $\$q$).

Monitor-extending Constructor. We now define a monitor-extending constructor \mathcal{E} that, given a monitored small-step semantics \rightarrow_{IF} , a partial function **Intercept** mapping statements to sequences of values, and an API register \mathcal{R}_{API} , produces a new monitored small-step semantics $\mathcal{E}(\rightarrow_{\text{IF}}, \text{Intercept}, \mathcal{R}_{\text{API}})$. The new extended semantics handles the invocation of APIs by applying the API relation that is returned by \mathcal{R}_{API} . API invocation is triggered by *interception points*, statements containing expression redexes (expressions that only have values as subexpressions) and that are in the domain of **Intercept**. Then, if the sequence of values to which its subexpressions evaluate is in the domain of the API register \mathcal{R}_{API} , their image by \mathcal{R}_{API} is the semantic relation that models the API to be executed.

The definition of \mathcal{E} , given in Figure 1, makes use of a syntactic function, **SubExpressions**, defined on JavaScript statements, such that **SubExpressions**[[s]] corresponds to the sequence comprising all the subexpressions of s in the order by which they are evaluated. Rules [NON-INTERCEPTED PROGRAM CONSTRUCT] and [INTERCEPTED PROGRAM CONSTRUCT - STANDARD EXECUTION] model the case in which the new small-step semantics behaves according to the original semantics \rightarrow_{IF} . Rule [INTERCEPTED PROGRAM CONSTRUCT - API EXECUTION] models the case in which an API is executed. The semantics rule retrieves the semantics relation that models the API to execute (using the API register) and then executes the API. After executing the API, the sequence of values of its subexpressions is replaced with the value to which the API call evaluates. Analogously, the sequence of levels of its subexpression is replaced with the reading effect of the API call.

3.3 Sufficient Conditions for Noninterferent API Extensions

We identify sufficient conditions to be satisfied by API relations in order for the new extended monitored semantics $\mathcal{E}(\rightarrow_{\text{IF}}, \text{Intercept}, \mathcal{R}_{\text{API}})$ to be noninterferent, assuming that the original monitor \rightarrow_{IF} is noninterferent.

The first condition requires that the API relation is *confined*, as formalized in Definition 5. An API relation is *confined* if it only creates/updates resources whose levels are higher than

$$\begin{array}{c}
 \text{NON-INTERCEPTED PROGRAM CONSTRUCT} \\
 \frac{s \notin \text{Intercept} \quad \langle \mu, s, \vec{\sigma}_{\text{pc}}, \Gamma, \Sigma, \vec{\sigma} \rangle \rightarrow_{\text{IF}} \langle \mu', s', \vec{\sigma}'_{\text{pc}}, \Gamma', \Sigma', \vec{\sigma}' \rangle}{\langle \mu, s, \vec{\sigma}_{\text{pc}}, \Gamma, \Sigma, \vec{\sigma} \rangle \mathcal{E}(\rightarrow_{\text{IF}}, \text{Intercept}, \mathcal{R}_{\text{API}}) \langle \mu', s', \vec{\sigma}'_{\text{pc}}, \Gamma', \Sigma', \vec{\sigma}' \rangle} \\
 \\
 \text{INTERCEPTED PROGRAM CONSTRUCT - STANDARD EXECUTION} \\
 \frac{s \in \text{Intercept} \quad (\mu, \text{SubExpressions}[[s]]) \notin \text{dom}(\mathcal{R}_{\text{API}}) \quad \langle \mu, s, \vec{\sigma}_{\text{pc}}, \Gamma, \Sigma, \vec{\sigma} \rangle \rightarrow_{\text{IF}} \langle \mu', s', \vec{\sigma}'_{\text{pc}}, \Gamma', \Sigma', \vec{\sigma}'_{\text{pc}} \rangle}{\langle \mu, s, \vec{\sigma}_{\text{pc}}, \Gamma, \Sigma, \vec{\sigma} \rangle \mathcal{E}(\rightarrow_{\text{IF}}, \text{Intercept}, \mathcal{R}_{\text{API}}) \langle \mu', s', \vec{\sigma}'_{\text{pc}}, \Gamma', \Sigma', \vec{\sigma}' \rangle} \\
 \\
 \text{INTERCEPTED PROGRAM CONSTRUCT - API EXECUTION} \\
 \frac{\begin{array}{c} s \in \text{Intercept} \quad (\mu, \text{SubExpressions}[[s]]) \in \text{dom}(\mathcal{R}_{\text{API}}) \\ |\text{SubExpressions}[[s]]| = n + 1 \quad \vec{\sigma} = \vec{\sigma}' \cdot \langle \sigma_0, \dots, \sigma_n \rangle \quad \Downarrow_{\text{API}} = \mathcal{R}_{\text{API}}(\mu, \text{SubExpressions}[[s]]) \\ \langle \mu, \Gamma, \Sigma, \text{Intercept}(s), \langle \sigma_0, \dots, \sigma_n \rangle \rangle \Downarrow_{\text{API}} \langle \mu', \Gamma', \Sigma', v', \sigma' \rangle \end{array}}{\langle \mu, s, \vec{\sigma}_{\text{pc}}, \Gamma, \Sigma, \vec{\sigma} \rangle \mathcal{E}(\rightarrow_{\text{IF}}, \text{Intercept}, \mathcal{R}_{\text{API}}) \langle \mu', v', \vec{\sigma}'_{\text{pc}}, \Gamma', \Sigma', \vec{\sigma}' \cdot \sigma' \rangle}
 \end{array}$$

■ **Figure 1** Definition of the monitor-extending constructor \mathcal{E} .

or equal to the least upper bound on the levels of its arguments. This constraint is needed because the choice of which API to execute may depend on all of its arguments.

► **Definition 5** (Confined API Relation/Register). An API relation \Downarrow_{API} is confined if, for every memory μ well-labeled by a labeling $\langle \Gamma, \Sigma \rangle$, every sequence of argument values \vec{v} and corresponding sequence of security levels $\vec{\sigma}$, if $\langle \mu, \langle \Gamma, \Sigma \rangle, \vec{v}, \vec{\sigma} \rangle \Downarrow_{\text{API}} \langle \mu', \langle \Gamma', \Sigma' \rangle, v', \sigma' \rangle$ for some memory μ' , labeling $\langle \Gamma', \Sigma' \rangle$, value v' , and level σ' ; then, for all security levels $\hat{\sigma}$:

$$\sqcup \vec{\sigma} \not\leq \hat{\sigma} \Rightarrow \mu, \Gamma, \Sigma \approx_{\text{id}, \hat{\sigma}} \mu', \Gamma', \Sigma' \wedge \sigma' \not\leq \hat{\sigma}$$

Furthermore, we say that the API Register function \mathcal{R}_{API} is confined, written $\mathbf{Conf}(\mathcal{R}_{\text{API}})$, if all the API relations in its range are confined, and if for every memories μ and μ' , labelings $\langle \Gamma, \Sigma \rangle$ and $\langle \Gamma', \Sigma' \rangle$, sequence of values \vec{v} , security level σ , and function β , such that $\mu, \Gamma, \Sigma \approx_{\beta, \sigma} \mu', \Gamma', \Sigma'$, then $\mathcal{R}_{\text{API}}(\mu, \vec{v}) = \mathcal{R}_{\text{API}}(\mu', \beta(\vec{v}))$.

The second condition requires that the API relation is *noninterferent*, as formalized in Definition 6. In order to relate the outputs of the API Register in two low-equal memories, we extend the notion of low-equality to API registers. Informally, two API registers are said to be low-equal if, whenever they are given as input two low-equal memories and two low-equal sequences of values, they output the same noninterferent API relation. Then, an API relation is *noninterferent* if whenever it is executed on two low-equal memories, it produces two low-equal memories and two low-equal values.

► **Definition 6** (Noninterferent API Relation/Register). An API relation \Downarrow_{API} is said to be noninterferent, written $\mathbf{NI}(\Downarrow_{\text{API}})$, if for every two memories μ_0 and μ_1 respectively well-labeled by $\langle \Gamma_0, \Sigma_0 \rangle$ and $\langle \Gamma_1, \Sigma_1 \rangle$, any two sequences of values \vec{v}_0 and \vec{v}_1 , respectively labeled by two sequences of security levels $\vec{\sigma}_0$ and $\vec{\sigma}_1$, and any security level σ for which there exists a function β such that $\vec{v}_0, \vec{\sigma}_0 \approx_{\beta, \sigma} \vec{v}_1, \vec{\sigma}_1$ and $\mu_0, \Gamma_0, \Sigma_0 \approx_{\beta, \sigma} \mu_1, \Gamma_1, \Sigma_1$, if:

$$\langle \mu_0, \Gamma_0, \Sigma_0, \vec{v}_0, \vec{\sigma}_0 \rangle \Downarrow_{\text{API}} \langle \mu'_0, \Gamma'_0, \Sigma'_0, v'_0, \sigma'_0 \rangle \wedge \langle \mu_1, \Gamma_1, \Sigma_1, \vec{v}_1, \vec{\sigma}_1 \rangle \Downarrow_{\text{API}} \langle \mu'_1, \Gamma'_1, \Sigma'_1, v'_1, \sigma'_1 \rangle$$

then there is an extension β' of β s.t. $\mu'_0, \Gamma'_0, \Sigma'_0 \approx_{\beta', \sigma} \mu'_1, \Gamma'_1, \Sigma'_1$ and $\langle v'_0 \rangle, \langle \sigma'_0 \rangle \approx_{\beta', \sigma} \langle v'_1 \rangle, \langle \sigma'_1 \rangle$.

Furthermore, we say that the API Register function \mathcal{R}_{API} is noninterferent, written $\mathbf{NI}(\mathcal{R}_{\text{API}})$, if all the API relations in its range are noninterferent and for every memories μ_0 and μ_1 , labelings $\langle \Gamma_0, \Sigma_0 \rangle$ and $\langle \Gamma_1, \Sigma_1 \rangle$, sequence of values \vec{v} , security level σ , and function β , such that $\mu_0, \Gamma_0, \Sigma_0 \approx_{\beta, \sigma} \mu_1, \Gamma_1, \Sigma_1$, then $\mathcal{R}_{\text{API}}(\mu_0, \vec{v}) = \mathcal{R}_{\text{API}}(\mu_1, \beta(\vec{v}))$.

► **Example 7** (Noninterferent JavaScript program using the Queue API). Assume that the APIs given in Example 1 are noninterferent and consider the following program that starts by computing two objects o_0 and o_1 :

```

1 q = queueAPI.createQueue();
2 if (h) { q.push(o1, 1); }
3 q.push(o0, 0); l = q.pop();

```

If this program starts with memories μ_i ($i \in \{0, 1\}$) using labeling $\langle \Gamma, \Sigma \rangle$ and assuming that in both executions the invocations of all the external APIs go through (i.e. the execution is never blocked), then it must terminate with memories μ'_i labeled by Γ', Σ :

$$\mu_i = \left[\begin{array}{l} (\#glob, o_0) \mapsto r_0, (\#glob, o_1) \mapsto r_1, \\ (\#glob, h) \mapsto i \end{array} \right] \quad \Gamma = \left[\begin{array}{l} (\#glob, h) \mapsto H, (\#glob, l) \mapsto L, \\ (\#glob, o_0) \mapsto L, (\#glob, o_1) \mapsto L \end{array} \right]$$

$$\mu'_i = \left[\begin{array}{l} (\#glob, o_0) \mapsto r_0, (\#glob, o_1) \mapsto r_1, \\ (\#glob, h) \mapsto i, (\#glob, l) \mapsto r_i, (\#glob, q) \mapsto r_q \end{array} \right] \quad \Gamma' = \left[\begin{array}{l} (\#glob, h) \mapsto H, (\#glob, l) \mapsto H, \\ (\#glob, o_0) \mapsto L, (\#glob, o_1) \mapsto L \end{array} \right]$$

Since initial memories are low-equal, $\mu_0, \Gamma, \Sigma \approx_{id, L} \mu_1, \Gamma, \Sigma$, we use the hypothesis that all three API relations are noninterferent to conclude that the memories yielded by the invocation of the API relations in lines 1, 2, and 3 are also low-equal. Furthermore, in the execution that maps h to 1, the value of l clearly depends on the value of h , from which we conclude that it is also the case in the execution that maps h to 0.

Our main result states that if the API relation is confined and noninterferent, then the extension of the noninterferent JavaScript monitor with the API monitor is noninterferent.

► **Theorem 8** (Security). *For every monitored semantics \rightarrow_{IF} , API register \mathcal{R}_{API} and set of interception points Intercept:*

$$\mathbf{NI}_{\text{mon}}(\rightarrow_{IF}) \wedge \mathbf{NI}(\mathcal{R}_{API}) \wedge \mathbf{Conf}(\mathcal{R}_{API}) \Rightarrow \mathbf{NI}_{\text{mon}}(\mathcal{E}(\rightarrow_{IF}, \text{Intercept}, \mathcal{R}_{API}))$$

4 A Meta-Compiler for Securing Web APIs

We now propose a way of extending an information flow monitor inlining compiler to take into account the execution of arbitrary APIs.

Input compilers. We assume available two inlining compilers specified by compilation functions \mathcal{C}_e and \mathcal{C}_s for compiling JavaScript expressions and statements, respectively. Function \mathcal{C}_s makes use of function \mathcal{C}_e . The compilers $\mathcal{C}_e/\mathcal{C}_s$ map every expression e /statement s to a pair $\langle s', i \rangle$, where:

1. statement s' simulates the execution of e/s in the monitored semantics;
2. index i is such that, after the execution of s' , (1) the compiler variable $\$ \hat{v}_i$ stores the value to which e/s evaluates in the original semantics and (2) the compiler variable $\$ \hat{l}_i$ stores its corresponding reading effect.

We assume that the inlining compiler works by pairing up each variable/property with a new one, called its *shadow* variable/property [10, 17], that holds its corresponding security level. Since the compiled program has to handle security levels, we include them in the set of program values, which means adding them to the syntax of the language as such, as well as adding two new binary operators corresponding to \leq (the order relation) and \sqcup (the least upper bound). Besides adding to every object o an additional shadow property $\$l_p$ for every property p in its domain, the inlined monitoring code is also assumed to extend o with a special property $\$struct$ that stores its structure security level.

► **Example 9** (Instrumented Labeling). Given an object $o = [p \mapsto v_0, q \mapsto v_1]$ pointed to by r_o and a labeling $\langle \Gamma, \Sigma \rangle$, such that $\Gamma(r_o) = [p \mapsto H, q \mapsto L]$ and $\Sigma(r_o) = L$, the instrumented counterpart of o labeled by $\langle \Gamma, \Sigma \rangle$ is $\hat{o} = [p \mapsto v_0, q \mapsto v_1, \$l_p \mapsto H, \$l_q \mapsto L, \$struct \mapsto L]$.

$$\begin{array}{c}
\text{INTERCEPTED EXPRESSION} \\
\text{SubExpressions}[[e]] = \langle e_0, \dots, e_n \rangle \quad \mathcal{C}_{\text{API}}\langle \mathcal{C}_e \rangle \langle e_0 \rangle = \langle s_0, i_0 \rangle \cdots \mathcal{C}_{\text{API}}\langle \mathcal{C}_e \rangle \langle e_n \rangle = \langle s_n, i_n \rangle \\
\hat{e} = \text{Replace}[[e, \$\hat{v}_{i_0}, \dots, \$\hat{v}_{i_n}]] \quad \langle \mathcal{C}_e \rangle \hat{e} = \langle \hat{s}, i \rangle \\
s_{\text{api}} = \left\{ \begin{array}{l} s_0 \dots s_n \\ \$if_{\text{sig}} = \$\text{apiRegister}(\$ \hat{v}_{i_0}, \dots, \$ \hat{v}_{i_n}); \\ \text{if}(\$if_{\text{sig}})\{ \\ \quad \$if_{\text{sig}}.\text{check}(\$ \hat{v}_{i_0}, \dots, \$ \hat{v}_{i_n}, \$ \hat{l}_{i_0}, \dots, \$ \hat{l}_{i_n}); \\ \quad \$\hat{v}_i = \hat{e}; \\ \quad \$ \hat{l}_i = \$if_{\text{sig}}.\text{label}(\$ \hat{v}_i, \$ \hat{v}_{i_0}, \dots, \$ \hat{v}_{i_n}, \$ \hat{l}_{i_0}, \dots, \$ \hat{l}_{i_n}); \\ \quad \} \text{ else } \{ \hat{s} \} \end{array} \right. \quad s' = \left\{ \begin{array}{ll} s_{\text{api}} & \text{if } \hat{e} \in \text{Intercept} \\ \hat{s} & \text{otherwise} \end{array} \right. \\
\hline
\mathcal{C}_{\text{API}}\langle \mathcal{C}_e \rangle \langle e \rangle = \langle s', i \rangle
\end{array}$$

■ **Figure 2** Extended Compiler \mathcal{C}_{API} .

4.1 IFlow Signatures

We propose *IFlow signatures* to simulate monitored executions of API relations. IFlow signatures are composed of three methods – *domain*, *check*, and *label*. Method *domain* checks whether or not to apply the API, *check* checks if the constraints associated with the API are verified, and *label* updates the instrumented labeling and outputs the reading effect associated with a call to the API. Functions *check* and *label* must be specified separately because *check* has to be executed before calling the API (in order to prevent its execution when it can potentially trigger a security violation), whereas *label* must be executed after calling the API (so that it can label the memory resulting from its execution). Formally, we define an *IFlow Signature* as a triple $\langle \#check, \#label, \#domain \rangle$, where: $\#check$ is the reference of the *check* function object, $\#label$ is the reference of the *label* function object, and $\#domain$ is the reference of the *domain* function object.

Runtime API Register. We assume the existence of a runtime function called the *runtime API register*, that simulates the API Register, which we denote by $\$apiRegister$. The function $\$apiRegister$ makes use of the *domain* method of each API in its range to decide whether there is an API relation associated with its inputs, in which case it outputs an object containing the corresponding IFlow Signature, otherwise it returns *null*.

Meta-compiler. Figure 2 presents a new meta-compiler, \mathcal{C}_{API} , that receives as input an inlining compiler for JavaScript expressions, \mathcal{C}_e , and outputs a new inlining compiler that can handle the invocation of the APIs whose signatures are in the range of the API register simulated by $\$apiRegister$. Since statement redexes are not intercepted, the compilation function \mathcal{C}_s is left unchanged except that it uses the new compilation function for expressions for compiling the subexpressions of the given statement. The specification of the meta-compiler makes use of a syntactic function Replace that receives as input an expression and a sequence of values and outputs the results from substituting each one of its subexpressions by the corresponding sequence value. **Intercept** is such that all the statements in its domain contain an expression redex whose execution is to be intercepted by the monitored semantics. Each expression that can be an interception point of the semantics is compiled by the compiler generated by the meta-compiler to a statement, which: (1) executes the statements corresponding to the compilation of its subexpressions, (2) tests if the sequence of values corresponding to the subexpressions of the expression to compile is associated with an IFlow signature, (3) if the test is true, it executes the *check* method of the corresponding IFlow signature, an expression equivalent to the original expression, and the *label* method of the corresponding IFlow signature. If the test is false, it executes the compilation of an

expression equivalent to the original one by the original inlining compiler. For simplicity, we do not take into account expressions that manipulate control flow, meaning that the evaluation of a given expression implies the evaluation of all its subexpressions. Therefore, we do not consider the JavaScript conditional expression. This limitation can be surpassed by re-writing all conditional expressions as IF statements before applying the compiler.

4.2 Correctness

We say that an inlining compiler is *correct* with respect to a given monitored semantics \rightarrow_{IF} if, provided that a program and its compiled counterpart are evaluated in “similar” memories, the evaluation of the original one in the monitored semantics terminates *if and only if* the evaluation of its compilation also terminates in the original semantics, in which case the final memories as well as the computed values are again “similar”. Here we use a notion of *similarity* between labeled memories in the monitored semantics and instrumented memories in the original semantics, denoted by \mathcal{S}_β . This relation requires that for every object in the labeled memory, the corresponding labeling coincides with the instrumented labeling and that the property values of the original object be similar to those of its instrumented counterpart. (The formal definition of \mathcal{S}_β can be found in the companion report [23].)

The correctness of the compiler generated by the meta-compiler depends on the correctness of the compiler given as input and the correctness of the IFlow signatures in the runtime API register. Definitions 10 and 11 formally specify the conditions that the instrumented API register must verify in order for the generated compiler to be correct. We use $\rightarrow_{\text{JS}}^*$ as the semantics relation for JavaScript configurations.

► **Definition 10 (Correct IFlow Signature).** An IFlow Signature $\langle \#c, \#l, \#d \rangle$ is *correct* with respect to an API \Downarrow_{API} if for all memories μ_0 and μ_1 , labeling $\langle \Gamma, \Sigma \rangle$, sequence of values \vec{v} , and sequence of security levels $\vec{\sigma}$, such that $\langle \mu_0, \Gamma, \Sigma \rangle \mathcal{S}_\beta \mu_1$ for some function β , then: $\langle \mu_0, \Gamma, \Sigma, \vec{v}, \vec{\sigma} \rangle \Downarrow_{\text{API}} \langle \mu'_0, \Gamma', \Sigma', v_0, \sigma \rangle$ if and only if (1) $\langle \mu_1, \#c(\beta(\vec{v})), \vec{\sigma} \rangle \rightarrow_{\text{JS}}^* \langle \mu'_1, \text{true} \rangle$, (2) $\langle \mu'_1, \beta(\vec{v}) \rangle \Downarrow_{\text{API}}^{\text{JS}} \langle \mu''_1, v_1 \rangle$, and (3) $\langle \mu'_1, \#l(v_1, \beta(\vec{v})), \vec{\sigma} \rangle \rightarrow_{\text{JS}}^* \langle \mu'''_1, \sigma \rangle$, in which case $\langle \mu'_0, \Gamma', \Sigma' \rangle \mathcal{S}_{\beta'} \mu'''_1$ and $v_0 \mathcal{S}_\beta v_1$, for some β' extending β .

► **Definition 11 (Correct Runtime API Register).** A runtime API register corresponding to a function object pointed by $\#\text{\$apiRegister}$ is correct with respect to an API register \mathcal{R}_{API} if for all memories μ_0 and μ_1 , labeling $\langle \Gamma, \Sigma \rangle$ and sequence of values \vec{v} , such that $\langle \mu_0, \Gamma, \Sigma \rangle \mathcal{S}_\beta \mu_1$ for some function β , then: $\mathcal{R}_{\text{API}}(\mu_0, \vec{v}) = \Downarrow_{\text{API}}$ if and only if (1) $\langle \mu_1, \#\text{\$apiRegister}(\beta(\vec{v})) \rangle \rightarrow_{\text{JS}}^* \langle \mu'_1, r_{\text{sig}} \rangle$, (2) $\langle \mu_0, \Gamma, \Sigma \rangle \mathcal{S}_{\beta'} \mu'_1$ for some β' extending β , and (3) signature $\langle o_{\text{sig}}(\text{“check”}), o_{\text{sig}}(\text{“label”}), o_{\text{sig}}(\text{“domain”}) \rangle$ is correct with respect to \Downarrow_{API} , where $o_{\text{sig}} = \mu'_1(r_{\text{sig}})$.

Theorem 12 states that provided that the compiler given as input to the meta-compiler is correct and the runtime API register is correct, the generated compiler is also correct.

► **Theorem 12 (Correctness).** *If compiler \mathcal{C} is correct w.r.t. \rightarrow_{IF} , then $\mathcal{C}_{\text{API}}(\mathcal{C})$ is correct w.r.t. $\mathcal{E}(\rightarrow_{\text{IF}}, \text{Intercept}, \mathcal{R}_{\text{API}})$ provided that the runtime API register is correct w.r.t. \mathcal{R}_{API} .*

5 Implementation of the Meta Compiler and DOM API Extension

An implementation of a meta-compiler based on the JavaScript inlined compiler of [24] can be found in [23] together with an online testing tool and a set of IFlow signatures that includes all those studied in the paper. As a case study, we describe the DOM API extension. Interaction between client-side JavaScript programs and the HTML document is done *via* the DOM API [19]. In order to access the functionalities of this API, JavaScript programs

$$\mathcal{R}_{\text{API}}^{\text{DOM}}(\mu, \langle r, m, \dots \rangle) = \begin{cases} \Downarrow_{\text{cre}} & \text{if } m = \text{"createElement"} \wedge r = \#doc \\ \Downarrow_{\text{app}} & \text{if } m = \text{"appendChild"} \wedge @tag \in \text{dom}(\mu(r)) \\ \Downarrow_{\text{rem}} & \text{if } m = \text{"removeChild"} \wedge @tag \in \text{dom}(\mu(r)) \\ \Downarrow_{\text{len}} & \text{if } m = \text{"length"} \wedge @tag \in \text{dom}(\mu(r)) \\ \Downarrow_{\text{par}} & \text{if } m = \text{"parentNode"} \wedge @tag \in \text{dom}(\mu(r)) \\ \Downarrow_{\text{ind}} & \text{if } m \in \text{Number} \wedge @tag \in \text{dom}(\mu(r)) \\ \Downarrow_{\text{sib}} & \text{if } m = \text{"nextSibling"} \wedge @tag \in \text{dom}(\mu(r)) \end{cases}$$

■ **Figure 3** API register $\mathcal{R}_{\text{API}}^{\text{DOM}}$ for the DOM API.

manipulate a special kind of objects, here named *DOM objects*. In contrast to the ECMA Standard [1] that specifies in full detail the internals of objects created during a JavaScript execution, the DOM API only specifies the behavior that DOM interfaces are supposed to exhibit when a program interacts with them. Hence, browser vendors are free to implement the DOM API as they see fit. In fact, in all major browsers, the DOM is not managed by the JavaScript engine. Instead, there is a separate engine, often called the *render engine*, whose role is to do so. Therefore, interactions between a JavaScript program and the DOM may potentially stop the execution of the JavaScript engine and trigger a call to the render engine. Thus, a monitored JavaScript engine has no access to the implementation of the DOM API. The special features of the DOM API spawn new implicit information flows [18, 21, 28].

► **Example 13** (Leak via removeChild - Order Leak). Suppose that in the original memory there are two orphan DIV nodes bound to variables `div1` and `div2`.

```
1  div1.appendChild(div2);
2  if(h) { div1.removeChild(div2); }
3  l = div1[0];
```

After the execution of this program, depending on the value of the high variable `h`, the value of the low variable `l` can be either that of `div2` or *undefined*, meaning that the final level associated with variable `l` must be *H* in both executions. This example shows that, when removing a node, the new indexes of its right siblings are affected. To tackle this problem, we specify the semantic relations corresponding to the DOM methods *removeChild* in such a way that for every DOM node the level of the property through which it is accessed is always lower than or equal to the levels of the properties corresponding to its right siblings.

As an example, we present the labeled API relation for removing a child with which we extend the JavaScript semantics for interaction with DOM objects. Function $\mathcal{R}_{\#Children}$ receives a memory μ as input and outputs a binary relation such that if $\langle r, n \rangle \in \mathcal{R}_{\#Children}(\mu)$, then the DOM node pointed to by r has n children (with indexes $0, \dots, n-1$). The API relation \Downarrow_{rem} removes the node pointed to by r_2 from the list of children of the node pointed to by r_0 , after checking that $\mu(r_0)$ is in fact the parent of $\mu(r_2)$. The object $\mu(r_0)$ is updated by shifting by -1 all the indexes equal to or higher than i (the index of the object being removed) and by removing index n . The levels of the indexes of the right siblings of the node to remove are accordingly shifted by -1 . The constraint of the rule prevents a program from removing in a high context a node that was inserted in a low context.

$$\begin{array}{c} \text{REMOVECHILD} \\ \mu(r_0)(i) = r_2 \quad \langle r_0, n+1 \rangle \in \mathcal{R}_{\#Children}(\mu) \quad \text{dom}(o_0) = \text{dom}(\gamma_0) = \text{dom}(\mu(r_0)) \setminus \{n\} \\ \forall_{0 \leq j < i} . o_0(j) = \mu(r_0)(j) \quad \forall_{i \leq j < n} . o_0(j) = \mu(r_0)(j+1) \quad o_0(@tag) = \mu(r_0)(@tag) \\ \forall_{0 \leq j < i} . \gamma_0(j) = \Gamma(r_0)(j) \quad \forall_{i \leq j < n} . \gamma_0(j) = \Gamma(r_0)(j+1) \quad \gamma_0(@tag) = \Gamma(r_0)(@tag) \\ \mu' = \mu[r_0 \mapsto o_0] \quad \Gamma' = \Gamma[r_0 \mapsto \gamma_0] \quad \sigma_0 \sqcup \sigma_1 \sqcup \sigma_2 \leq \Gamma(r_0)(i) \\ \hline \langle \mu, \Gamma, \Sigma, \langle r_0, _, r_2 \rangle, \langle \sigma_0, \sigma_1, \sigma_2 \rangle \rangle \Downarrow_{\text{rem}} \langle \mu', \Gamma', \Sigma, r_2, \sigma_0 \sqcup \sigma_1 \sqcup \sigma_2 \rangle \end{array}$$

```

domain = function(o0, m){return o0[@tag] &&& (m == "removeChild"); }

check = function(o0, m, o2, σ0, σ1, σ2){
  var i = 0;
  while(i ≤ o0["length"]){
    if(o0[i] == o2) break;
  }
  $check(i < o0["length"]);
  o2[$index] = i;
  return $check(σ0 ⊔ σ1 ⊔ σ2 ≤ o0[$shadow(i)]);
}

label = function(ret, o0, m, o2, σ0, σ1, σ2){
  var i = o2[$index];
  var j = o0["length"];
  while(j > i){
    o0[$shadow(j)] = o0[$shadow(j + 1)];
  }
  delete o0[$shadow(o0["length"])];
  delete o2[$index];
  return σ0 ⊔ σ1 ⊔ σ2;
}

```

■ **Figure 4** IFlow Signature of \Downarrow_{rem} .

In order for these API relations to be added to the semantics, one has to add them to the API register. Hence, we assume that the \mathcal{R}_{API} extends the API register given in Figure 3.

The following two lemmas validate the hypotheses of the security theorem (Theorem 8) for $\mathcal{R}_{\text{API}}^{\text{DOM}}$, allowing us to conclude that the extension of a noninterferent JavaScript monitor with the DOM API is noninterferent.

► **Lemma 14** (Confinement for the DOM API). $\text{Conf}(\mathcal{R}_{\text{API}}^{\text{DOM}})$.

► **Lemma 15** (Noninterference for the DOM API). $\text{NI}(\mathcal{R}_{\text{API}}^{\text{DOM}})$.

Figure 4 presents a possible IFlow signature for the API relation \Downarrow_{rem} , which makes use of the two following runtime functions: (1) $\$check$ diverges if its argument is different from *true* and returns *true* otherwise and (2) $\$shadow$ receives as input a property name and outputs the name of the corresponding shadow property.

6 Conclusion

In summary, we have proposed a methodology for extending arbitrary monitored JavaScript semantics with secure APIs, which allows to prove the security of the extended monitor in a modular way. As a case study, we extend the inlined compiler of [24] with a fragment of the DOM Core Level 1 API. The proofs of the results as well as an implementation that includes the IFlow signatures of the APIs studied in the paper can be found in [23].

References

- 1 The 5.1th edition of ECMA 262 June 2011. ECMAScript Language Specification. Technical report, ECMA, 2011.
- 2 T. H. Austin and C. Flanagan. Efficient purely-dynamic information flow analysis. In *PLAS*, 2009.
- 3 T. H. Austin and C. Flanagan. Permissive dynamic information flow analysis. In *PLAS*, 2010.
- 4 T. H. Austin and C. Flanagan. Multiple facets for dynamic information flow. In *POPL*, 2012.
- 5 A. Banerjee and D. A. Naumann. Secure information flow and pointer confinement in a java-like language. In *CSFW*, 2002.

- 6 N. Bielova. Survey on javascript security policies and their enforcement mechanisms in a web browser. *Special Issue on Automated Specification and Verification of Web Systems of JLAP*, 2013. To appear.
- 7 Nataliia Bielova, Dominique Devriese, Fabio Massacci, and Frank Piessens. Reactive Non-interference for a Browser Model. In *5th International Conference on Network and System Security*, 2011.
- 8 A. Birgisson, D. Hedin, and A. Sabelfeld. Boosting the permissiveness of dynamic information-flow tracking by testing. In *ESORICS*, 2012.
- 9 A. Bohannon, B. C. Pierce, V. Sjöberg, S. Weirich, and S. Zdancewic. Reactive noninterference. In *ACM Conference on Computer and Communications Security*, 2009.
- 10 A. Chudnov and D. A. Naumann. Information flow monitor inlining. In *CSF*, 2010.
- 11 D. E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5), 1976.
- 12 D. Devriese and F. Piessens. Noninterference through secure multi-execution. In *S&P*, 2010.
- 13 P. Gardner, G. Smith, M. J. Wheelhouse, and U. Zarfaty. Dom: Towards a formal specification. In *PLAN-X*, 2008.
- 14 G. Le Guernic. *Confidentiality Enforcement Using Dynamic Information Flow Analyses*. PhD thesis, Kansas State University, 2007.
- 15 A. Guha, B. Lerner, J. Gibbs Politz, and S. Krishnamurthi. Web api verification: Results and challenges. 2012.
- 16 D. Hedin and A. Sabelfeld. Information-flow security for a core of javascript. In *CSF*, 2012.
- 17 J. Magazinius, A. Russo, and A. Sabelfeld. On-the-fly inlining of dynamic security monitors. *Computers & Security*, 2012.
- 18 Ana Almeida Matos, José Fragoso Santos, and Tamara Rezk. An Information Flow Monitor for a Core of DOM - Introducing References and Live Primitives. In *Trustworthy Global Computing - 9th International Symposium (TGC), Rome, Italy, 2014*.
- 19 W3C Recommendation. DOM: Document Object Model (DOM). Technical report, W3C, 2005.
- 20 A. Russo and A. Sabelfeld. Dynamic vs. static flow-sensitive security analysis. In *CSF*, 2010.
- 21 A. Russo, A. Sabelfeld, and A. Chudnov. Tracking information flow in dynamic tree structures. In *ESORICS*, 2009.
- 22 A. Sabelfeld and A. C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 2003.
- 23 José Fragoso Santos and Tamara Rezk. Information flow monitor-inlining compiler. <http://j3fsantos.github.io/PersonalPage/IFMonitor/>.
- 24 José Fragoso Santos and Tamara Rezk. An Information Flow Monitor-Inlining Compiler for Securing a Core of Javascript. In *ICT Systems Security and Privacy Protection - 29th IFIP TC 11 International Conference (SEC), Marrakech, Morocco, 2014*.
- 25 P. Shroff, S. F. Smith, and M. Thober. Dynamic dependency monitoring to secure information flow. In *CSF*, 2007.
- 26 A. Taly, U. Erlingsson, J. C. Mitchell, M. S. Miller, and J. Nagra. Automated analysis of security-critical javascript apis. In *SP*, 2011.
- 27 V. N. Venkatakishnan, W. Xu, D. C. DuVarney, and R. Sekar. Provably correct runtime enforcement of non-interference properties. In *ICICS*, 2006.
- 28 Deepak Garg Vineet Rajani, Abhishek Bichhawat and Christian Hammer. Information Flow control for Event Handling and the DOM in Web Browsers. In *CSF*, 2015. to appear.