Hybrid Typing of Secure Information Flow in a JavaScript-Like Language

José Fragoso Santos[1(✉)], Thomas Jensen[2], Tamara Rezk[2], and Alan Schmitt[2]

[1] Imperial College London, London, UK
jose.fragoso.santos@imperial.ac.uk
[2] Inria, Sophia Antipolis, France
{thomas.jensen,tamara.rezk,alan.schmitt}@inria.fr

Abstract. As JavaScript is highly dynamic by nature, static information flow analyses are often too coarse to deal with the dynamic constructs of the language. To cope with this challenge, we present and prove the soundness of a new hybrid typing analysis for securing information flow in a JavaScript-like language. Our analysis combines static and dynamic typing in order to avoid rejecting programs due to imprecise typing information. Program regions that cannot be precisely typed at static time are wrapped inside an internal *boundary* statement used by the semantics to interleave the execution of statically verified code with the execution of code that must be dynamically checked.

1 Introduction

The dynamic aspects of JavaScript make the analysis of JavaScript programs very challenging. On one hand, one may use a purely static analysis, but either restrict the language to exclude these dynamic aspects or over-approximate them; this is too coarse to be applicable in practice. On the other hand, one may use purely dynamic mechanisms, such as monitoring or secure multi-executions [1,6,8,16]; but the gained precision comes at the cost of a much lower performance compared to the original code [7].

We propose a general hybrid analysis to statically verify secure information flow in a core of JavaScript. Following the hybrid typing motto "static analysis where possible with dynamic checks where necessary" [5], we are able to reduce the runtime overhead introduced by purely dynamic analyses without excluding dynamic field operations. In fact, our analysis can handle some of the most challenging JavaScript features, such as prototype-based inheritance, extensible objects, and constructs for checking the existence of object properties. Its key ingredient is an internal boundary statement inspired by recent work in inter-language interoperability [10]. The static component of our analysis wraps program regions that cannot be precisely verified inside an internal *boundary* statement instead of rejecting the whole program. This boundary statement identifies the regions of the program that must be verified at runtime—which may be as small as a single statement—and enables the initial set up required

© Springer International Publishing Switzerland 2016
P. Ganty and M. Loreti (Eds.): TGC 2015, LNCS 9533, pp. 63–78, 2016.
DOI: 10.1007/978-3-319-28766-9_5

by the dynamic analysis. In summary, the proposed boundary statement allows the semantics to effortlessly interleave the execution of statically verified code with the execution of code that must be verified at runtime.

Although our work is generally motivated by the verification of dynamic features of JavaScript, we choose to focus on the particular case of constructs that rely on dynamic computation of object field names, which we call *dynamic field operations*. In JavaScript, one can access a field `f` of an object `o` either by writing `o.f` or `o[e]`, where `e` is an expression that dynamically evaluates to the string `f`. Dynamic computation of field names is one of the major sources of imprecision of static analyses for JavaScript [9].

Example 1 (Running example: the challenge of typing dynamic field operations). Below we present a program that creates an object `o` with a secret field `secret1` and two public fields `public1` and `public2`.

```
o = {}; o.secret1 = secret_input();
o.public1 = public_input(); o.public2 = public_input(); public = o[g()]
```

The secret field `secret1` gets a secret input via function `secret_input`, while the two public fields `public1` and `public2` each get a public input via function `public_input`. The program then assigns the value of one of the three fields to the public variable `public`, as determined by the return value of function `g`. Concretely, when `g` returns the string `"secret1"`, the program assigns a secret value to `public` and the execution is insecure. On the other hand, when `g` returns either `"public1"` or `"public2"`, the program assigns a public value to `public` and the execution is secure. However, in order to make sure that `g` never returns `"secret1"`, a static analysis needs to predict the dynamic behaviour of `g`, which is, in general, undecidable.

The loss of precision introduced by the dynamic computation of field names is not exclusive to field projections. It also occurs in method calls, field deletions, and membership checks. We account for the use of these operations by verifying them at runtime. When verifying a statement containing a dynamic field operation, the static component of the analysis wraps it inside a boundary statement. In the case of the running example, all statements except the last one are statically typed. In contrast, the last assignment is re-written as `@monitor(@type_env, @pc, @ret, public = o[g()])`, where the first three arguments of the monitor statement are used for the setup of the runtime analysis. Hence, when the program is executed the only overhead introduced by the dynamic component of our hybrid analysis regards the security checks for validating or rejecting the statement `public = o[g()]`.

Contributions. The main contribution of the paper is the design of a new hybrid analysis for verifying secure information flow in a JavaScript-like language. To achieve this, we introduce: **(1)** a type language specifically designed to control information flow in a subset of JavaScript, **(2)** a static type system for verifying statements not containing dynamic field operations, **(3)** a dynamic typing analysis for verifying statements containing dynamic field operations, and **(4)** a novel boundary operator for interleaving the execution of statically verified

Table 1. Core JS syntax - expressions and statements

$v \in \mathcal{V}al$	$::=$	$lit \mid \underline{lit} \mid l \mid \lambda x : \dot{\tau}.s$
$e \in \mathcal{E}xpr$	$::=$	$v \mid \mathsf{this} \mid x \mid x = e \mid \{\,\}\,[\dot{\tau}] \mid e.f \mid e_1[e_2] \mid e_1.f = e_2$
		$\mid e_1[e_2] = e_3 \mid f \,\mathsf{in}\, e \mid [e_1] \,\mathsf{in}\, e_2 \mid \mathsf{delete}\, e.f \mid \mathsf{delete}\, e_1[e_2]$
		$\mid \mathsf{function}\,(x)[\dot{\tau}]\{s\} \mid e_1(e_2) \mid e_1.x(e) \mid e_1[e_2](e_3)$
$s \in \mathcal{S}tmt$	$::=$	$e \mid \mathsf{var}\, x\, [\dot{\tau}] \mid s_1;\, s_2 \mid \mathsf{if}(e)\,\{s_1\}\,\mathsf{else}\,\{s_2\} \mid \mathsf{return}\, e$

regions with dynamically verified ones. Finally, we have implemented a prototype as well as a case study, available online at [15].

2 Core JS

Syntax. The syntax of Core JS is given in Table 1. Expressions include values, the keyword this, variables, variable assignments, object literals, static and dynamic field projections, static and dynamic field assignments, static and dynamic membership checks, static and dynamic field deletions, function literals, function calls, and static and dynamic method calls. Statements include expression statements, variable declarations, sequences, conditional statements, and return statements. We distinguish two types of *values*: literal values and runtime values. Literal values include numbers, booleans, strings, and undefined. Runtime values, ranged over by $\underline{v}$, include parsed literal values, locations, and parsed function literals. Object literals, function literals, and variable declarations are annotated with their respective *security types* (which are explained in Sect. 3). In the following, we use $\mathcal{E}xpr_{\not{\,}}$ for the set of Core JS dynamic field operations.

Memory Model. A heap $H \in \mathcal{H}eap : \mathcal{L}oc \times \mathcal{X} \rightharpoonup \mathcal{V}al$ is a partial mapping from locations in $\mathcal{L}oc$ and field names in $\mathcal{X}$ to values in $\mathcal{V}al$. We denote a heap cell by $(l, f) \mapsto v$, the union of two disjoint heaps by $H_1 \uplus H_2$, a read operation by $H(l, f)$, and a heap update operation by $H[l.f \mapsto v]$. An object can be seen as a set of heaps cells addressed by the same location but with different field names. We use $l \mapsto \{f_1 : v_1, \ldots, f_n : v_n\}$ as an abbreviation for the object $(l, f_1) \mapsto v_1 \uplus \ldots \uplus (l, f_n) \mapsto v_n$.

Every object has a prototype, whose location is stored in a special field *_proto_*. In order to determine the value of a field f of an object o, the semantics first checks whether f is one of the fields of o. If that is the case, the field look-up yields that value. Otherwise, the semantics checks whether f belongs to the fields of the prototype of o and so forth. The sequence of objects that can be accessed from a given object through the inspection of the respective prototypes is called a prototype chain. The prototype chain inspection procedure is modelled by the semantic function π given in appendix. Informally, the expression $\pi(H, l, f)$ denotes the location of the first object that defines f in the prototype chain of the object pointed to by l (if no such object exists, π returns null). Given that most

Table 2. Evaluation contexts

$$\hat{E} ::= \square \mid x = \hat{E} \mid \hat{E}.f \mid \hat{E}[e] \mid l[\hat{E}] \mid \hat{E}.f = e \mid \hat{E}[e_1] = e_2$$
$$\mid l[\hat{E}] = e \mid l[f] = \hat{E} \mid [\hat{E}]\ \mathsf{in}\ e \mid [f]\ \mathsf{in}\ \hat{E} \mid \mathsf{delete}\ \hat{E}.f \mid \mathsf{delete}\ \hat{E}[e]$$
$$\mid \mathsf{delete}\ l[\hat{E}] \mid \hat{E}(e) \mid l(\hat{E}) \mid \hat{E}.f(e) \mid \hat{E}[e](e) \mid l[\hat{E}](e) \mid l[f](\hat{E})$$
$$E ::= \hat{E} \mid E;\ s \mid \mathsf{if}(\hat{E})\ \{s_1\}\ \mathsf{else}\ \{s_2\} \mid \mathsf{return}\ \hat{E}$$

implementations of JavaScript allow for explicit prototype mutation, we include this feature in Core JS. For instance, $x._proto_$ evaluates to the prototype of the object bound to x and $x._proto_ = y$ sets the prototype of the object bound to x to the object bound to y.

Scope is modelled using *environment records*. An environment record is simply an internal object that maps variable names to their respective values. An environment record is created for every function or method call. We use $\mathsf{act}(l, x, v, s, l')$ to denote the environment record that: **(1)** is identified by location l where it is stored, **(2)** maps x to v, **(3)** maps all the variables declared in s to $\mathsf{undefined}$, and **(4)** maps the field @*this* to the location l'. (Note that environment records map a single variable because functions have a single argument. Moreover, in the execution of a method call, the field @*this* is used to store the location of the object on which the method was invoked.) Variables are resolved with respect to a list of environment record locations, called *scope chain*. The variable inspection procedure is modelled by the semantic function σ given in appendix. We let $\sigma(H, L, x)$ denote the location of the first environment record that defines x in the scope chain L. The global object, assumed to be pointed to by a fixed location l_g, is the environment record that binds global variables.

Since functions are first-class citizens, the evaluation of a function literal triggers the creation of a special type of object, called *function object*. Every function object has two fields: @*body* and @*scope*, which respectively store the corresponding parsed function literal and the scope chain that was active when the function literal was evaluated. Functions execute in the scope in which the they were evaluated.

Semantics. Figure 1 presents a fragment of the semantics of Core JS in the style of Wright and Felleisen [19] (the full semantics is given in appendix). A configuration Ψ has the form $\langle H, L, s\rangle$ where H is the current heap, L the current scope chain, and s the statement to execute. Transitions are labelled with an internal event α for the use of the dynamic analysis. The evaluation order is specified with the help of evaluation contexts, whose syntax is given in Table 2. In the following, we use $l :: L$ for the list obtained by prepending l to L and $\mathsf{head}(L)$ for the first element of L.

Rule VARIABLE uses σ to determine the location l' of the environment record that defines x and reads its value from the heap. Rule DYN FIELD PROJECTION uses π to determine the location l'' of the object that defines f in the prototype chain of the object pointed to by l' and then reads its value from the heap. Rule

VARIABLE
$$\frac{l = \mathsf{head}(L) \qquad l' = \sigma(H, L, x) \qquad \underline{v} = H(l', x)}{\langle H, L, x\rangle \overset{\mathsf{var}_l(x)}{\rightarrow} \langle H, L, \underline{v}\rangle}$$

DYN. FIELD PROJECTION
$$\frac{l = \mathsf{head}(L) \qquad l'' = \pi(H, l', f) \qquad \underline{v} = H(l'', f)}{\langle H, L, l'[f]\rangle \overset{\mathsf{f\text{-}proj}_l(f)}{\rightarrow} \langle H, L, \underline{v}\rangle}$$

DYN. FIELD ASSIGNMENT
$$\frac{l' = \mathsf{head}(L) \qquad H' = H[l.f \mapsto \underline{v}]}{\langle H, L, l[f] = \underline{v}\rangle \overset{\mathsf{f\text{-}ass}_{l'}(f)}{\rightarrow} \langle H', L, \underline{v}\rangle}$$

MEMBERSHIP CHECK - TRUE
$$\frac{l' = \mathsf{head}(L) \qquad \pi(H, l, f) \neq \mathsf{null}}{\langle H, L, [f]\ \mathsf{in}\ l\rangle \overset{\mathsf{in}_{l'}(f)}{\rightarrow} \langle H, L, \mathsf{true}\rangle}$$

FUNCTION LITERAL
$$\frac{l = \mathsf{head}(L) \qquad l' = \mathsf{fresh}(H, \dot{\tau}) \qquad H' = H \uplus l' \mapsto \{@scope : L, @body : \lambda x : \dot{\tau}.s\}}{\langle H, L, \mathsf{function}\ (x)[\dot{\tau}]\{s\}\rangle \overset{\mathsf{push}_l(\dot{\tau})}{\rightarrow} \langle H', L, l'\rangle}$$

FUNCTION CALL
$$\frac{l' = \mathsf{head}(L) \quad l'' \notin \mathsf{dom}(H) \quad \lambda x : \dot{\tau}.s = H(l, @body) \quad L' = H(l, @scope) \quad H' = H \uplus \mathsf{act}(l'', x, \underline{v}, s, l_g)}{\langle H, L, l(\underline{v})\rangle \overset{\mathsf{f\text{-}call}_{l'}}{\rightarrow} \langle H', l'' :: L', @\mathsf{FunExe}(L, s)\rangle}$$

IF END
$$\frac{l = \mathsf{head}(L)}{\langle H, L, @\mathsf{EI}(\underline{v})\rangle \overset{\succ l}{\rightarrow} \langle H, L, \underline{v}\rangle}$$

IF - TRUE
$$\frac{l = \mathsf{head}(L) \quad \neg\mathsf{false}(\underline{v}) \quad s' = @\mathsf{EI}(s_1)}{\langle H, L, \mathsf{if}(\underline{v})\ \{s_1\}\ \mathsf{else}\ \{s_2\}\rangle \overset{\mathsf{if}\,l}{\rightarrow} \langle H, L, s'\rangle}$$

CONTEXTUAL PROPAGATION
$$\frac{\langle H, L, s\rangle \overset{\alpha}{\rightarrow} \langle H', L', s'\rangle}{\langle H, L, E[s]\rangle \overset{\alpha}{\rightarrow} \langle H', L', E[s']\rangle}$$

Fig. 1. Fragment of the small-step semantics of core JS

DYN FIELD ASSIGNMENT updates the current heap with a mapping from l and f to $\underline{v}$. Rule MEMBERSHIP CHECK - TRUE checks if f is defined in the prototype chain of the object pointed to by l and evaluates to true. Rule FUNCTION LITERAL adds a new function object to the heap. Rule FUNCTION CALL extends the heap with a new environment record for the evaluation of the function pointed to by l. The current scope chain L is replaced with the scope chain L' that was active when the corresponding function literal was evaluated extended with the location l'' of the newly created environment record. The semantics makes use of an internal statement $@\mathsf{FunExe}(L, s)$ for keeping track of the caller's scope chain during the execution of the function's body. Rule IF - TRUE checks if the guard of the conditional does not belong to the set of *falsy* values $-\{\mathsf{false}, 0, \mathsf{undefined}, \mathsf{null}\}-$ and replaces the whole conditional with its then-branch followed by an internal statement @EI for notifying the dynamic analysis of the end of that branch. CONTEXTUAL PROPAGATION is standard.

3 Static Typing Secure Information Flow in Core JS

In this section, we present both a new type language for controlling information flow in JavaScript and the static component of our analysis. Here, the specification of security policies relies on two key elements: a lattice of security levels and a

typing environment that maps resources to security types, which can be viewed as safety types annotated with security levels. In the examples, we use $\mathcal{L} = \{H, L\}$ with $L \sqsubset H$, meaning that L-labelled resources (*low* resources) are less confidential than those labelled with H (*high*). We use $\sqcup$, $\bot$, and $\top$ for the least upper bound (*lub*), the *bottom* level, and the *top* level, respectively.

Security Types. A security type $\dot{\tau} = \tau^\sigma$ is obtained by pairing up a *raw type* τ with a security level σ, called its *external level.* The external level of a security type establishes an upper bound on the levels of the resources on which the values of that type may depend. For instance, a primitive value of type $\textsf{PRIM}^L$ may only depend on *low* resources. The syntax of raw types is given and explained below:

$$\begin{aligned}\tau ::= \textsf{PRIM} \mid \ \langle \dot{\tau}.\dot{\tau} \xrightarrow{\sigma} \dot{\tau} \rangle \mid \ \langle \kappa.\dot{\tau} \xrightarrow{\sigma} \dot{\tau} \rangle \\ \mid \mu\kappa.\langle f^\sigma : \dot{\tau}, \cdots, f^\sigma : \dot{\tau}, *^\sigma : \dot{\tau} \rangle \mid \mu\kappa.\langle f^\sigma : \dot{\tau}, \cdots, f^\sigma : \dot{\tau} \rangle\end{aligned}$$

- The type $\textsf{PRIM}$ is the type of expressions which evaluate to primitive values.
- The type $\langle \dot{\tau}_0.\dot{\tau}_1 \xrightarrow{\sigma} \dot{\tau}_2 \rangle$ is the type of expressions which evaluate to functions that map values of type $\dot{\tau}_1$ to values of type $\dot{\tau}_2$ and during the execution of which, the keyword $\textsf{this}$ is bound to an object of type $\dot{\tau}_0$. Level σ is the *writing effect* [14] of functions of this type, that is, a lower bound on the levels of the resources updated or created during their execution. When specifying a function type inside an object type, one can use the type variable bound by that object type as the type of the keyword $\textsf{this}$ (in the syntax of types, κ ranges over the set of type variables).
- The type $\mu\kappa.\langle f_0^{\sigma_0} : \dot{\tau}_0, \cdots, f_n^{\sigma_n} : \dot{\tau}_n, *^{\sigma_*} : \dot{\tau}_* \rangle$ is the type of expressions which evaluate to objects that *may* define the fields f_0 to f_n mapping each field f_i to a value of security type $\dot{\tau}_i$. The security type assigned to $*$ is the *default security type*, which is the security type of all fields not in $\{f_0, \cdots, f_n\}$. Every field f_i is further associated with an *existence level* σ_i that establishes an upper bound on the levels of the contexts in which the field can be created or deleted. The level σ_* is the *default existence level.* When no default security type is declared, the objects of the type may only define explicitly declared fields.

  The reason why we do not precisely track the presence of fields in object types is that we do not want the type of an object to change at runtime even though its structure may change. Notice that the absence of a field in a type does not mean it cannot be accessed in objects of that type: this field may still be defined in the prototype chain. We could have flattened security types for objects by requiring every object type to explicitly declare all the fields accessible through the prototype chains of the objects of that type, but this would have two disadvantages. First, object types would be less precise, and second, they would be much larger as the types of prototype fields would be duplicated. The cost of this design choice is a more complex Static Field Projection typing rule that has to take the prototype chain into account.

Given a security type $\dot{\tau}$, the expression $\textsf{lev}(\dot{\tau})$ denotes its external level and $\lfloor \dot{\tau} \rfloor$ its raw type (for instance, $\textsf{lev}(\textsf{PRIM}^L) = L$ and $\lfloor \textsf{PRIM}^L \rfloor = \textsf{PRIM}$). We define $\dot{\tau}^\sigma$

Table 3. Typing environment for the Examples 1 to 6

$$\Gamma(\texttt{public}) = \textsf{PRIM}^L \qquad \Gamma(\texttt{secret}) = \textsf{PRIM}^H$$
$$\Gamma(\texttt{secret_input}) = \langle \dot\tau_g._ \xrightarrow{H} \textsf{PRIM}^H\rangle^L \qquad \Gamma(\texttt{public_input}) = \langle \dot\tau_g._ \xrightarrow{H} \textsf{PRIM}^L\rangle^L$$
$$\dot\tau_o = \mu\kappa\ .\left\langle \begin{array}{l}\texttt{public1}^L : \textsf{PRIM}^L,\\ \texttt{public2}^L : \textsf{PRIM}^L,\\ \texttt{secret1}^H : \textsf{PRIM}^H\\ \texttt{secret2}^H : \textsf{PRIM}^H\end{array}\right\rangle^L$$
$$\Gamma(\texttt{g}) = \langle \dot\tau_g._ \xrightarrow{H} \textsf{PRIM}^L\rangle^L \quad \Gamma(\texttt{o0}) = \mu\kappa.\langle \texttt{_proto_}^H : \dot\tau_o\rangle^L$$
$$\Gamma(\texttt{o}) = \Gamma(\texttt{o1}) = \Gamma(\texttt{o2}) = \dot\tau_o$$

as $\lfloor\dot\tau\rfloor^{\mathsf{lev}(\dot\tau)\sqcup\sigma}$ (for example, $(\textsf{PRIM}^L)^H = \textsf{PRIM}^H$). Given a function security type $\dot\tau = \langle\dot\tau_0.\dot\tau_1 \xrightarrow{\sigma} \dot\tau_2\rangle^{\sigma'}$, we use $\dot\tau.\mathsf{this}$, $\dot\tau.\mathsf{arg}$, $\dot\tau.\mathsf{ret}$, and $\dot\tau.\mathsf{wef}$ to denote $\dot\tau_0$, $\dot\tau_1$, $\dot\tau_2$, and σ, respectively. Given an object security type $\dot\tau$, we use $\mathsf{dom}(\dot\tau)$ for the set containing all field names explicitly declared in $\dot\tau$ (including $*$, if present). Given a field name f and an object security type $\dot\tau$, $\dot\tau.f$ ($\dot\tau.\overline{f}$, resp.) denotes either the security type (existence level resp.) with which $\dot\tau$ associates f or its default security type (existence level, resp.) when $f \notin \mathsf{dom}(\dot\tau)$ and $* \in \mathsf{dom}(\dot\tau)$. The ordering $\sqsubseteq$ on security levels induces a simple ordering $\preceq$ on security types: $\dot\tau_0 \preceq \dot\tau_1$ *iff* $\mathsf{lev}(\dot\tau_0) \sqsubseteq \mathsf{lev}(\dot\tau_1)$ and $\lfloor\dot\tau_0\rfloor = \lfloor\dot\tau_1\rfloor$. We use $\dot\tau_g$ for the type of the global object. Finally, a typing environment Γ is simply a mapping from variables to security types.

Example 2. Table 3 presents the typing environment used to type the programs given in Examples 1 to 6. Since `secret_input`, `public_input`, and `g` are to be used as functions, their respective types use the type of the global object as the type of the keyword this. Since none of these three functions expects an argument or updates the heap, their respective types omit the type of the argument and declare a *high* writing effect. Our design choice of not flattening object types can also be seen in this example: the type of `o0` is much shorter as it does not need to mention at top level the fields declared in $\dot\tau_o$.

Static Type System. The key insight of the static type system is that it wraps program regions which cannot be precisely analysed at static time within a boundary statement $@monitor(\Gamma, pc, \dot\tau_r, s)$ responsible for turning on the typing analysis at runtime. The parameters Γ, pc, and $\dot\tau_r$ are the typing environment, the context level [14], and the type of the function whose body is being typed, respectively. Given a typing environment Γ, a level pc, and an expression e, the typing judgment $\Gamma, pc \vdash_e e \hookrightarrow e' : \dot\tau$ means that e is rewritten as a semantically equivalent expression e', which may include boundary statements, has raw type $\lfloor\dot\tau\rfloor$, and reads variables or fields of level at most $\mathsf{lev}(\dot\tau)$. Typing judgements for statements, with the form $\Gamma, pc, \dot\tau_r \vdash_s s \hookrightarrow s'$, differ from typing judgements for expressions in that they do not assign a type to the statement. When e (s resp.) coincides with e' (s' resp.), we omit $\hookrightarrow e'$ ($\hookrightarrow s'$ resp.) from the typing rules. The most relevant typing rules are given in Fig. 2 and described below. (We omit

$$\textsc{Literal}\quad \Gamma, pc \vdash_e lit : \mathsf{PRIM}^{\bot}$$

$$\textsc{Variable}\quad \Gamma, pc \vdash_e x : \Gamma(x)$$

$$\textsc{Assignment}\quad \frac{\Gamma, pc \vdash_e e : \dot{\tau} \qquad \dot{\tau}^{pc} \preceq \Gamma(x)}{\Gamma, pc \vdash_e x = e : \dot{\tau}}$$

$$\textsc{Static Field Projection}\quad \frac{\Gamma, pc \vdash_e e : \dot{\tau} \qquad \dot{\tau}_f = \pi(\dot{\tau}, f)}{\Gamma, pc \vdash_e e.f : \dot{\tau}_f^{\mathsf{lev}(\dot{\tau})}}$$

$$\textsc{Static Member Check}\quad \frac{\Gamma, pc \vdash_e e : \dot{\tau} \qquad \sigma = \mathsf{lev}(\dot{\tau}) \sqcup \overline{\pi}(\dot{\tau}, f)}{\Gamma, pc \vdash_e f \,\mathsf{in}\, e : \mathsf{PRIM}^{\sigma}}$$

$$\textsc{Static Field Assignment}\quad \frac{\forall_{i=1,2}\, \Gamma, pc \vdash_e e_i : \dot{\tau}_i \qquad \dot{\tau}_2 \preceq \dot{\tau}_1.f \quad pc \sqcup \mathsf{lev}(\dot{\tau}_1) \sqsubseteq \dot{\tau}_1.\overline{f}}{\Gamma, pc \vdash_e e_1.f = e_2 : \dot{\tau}_2}$$

$$\textsc{Static Field Deletion}\quad \frac{\Gamma, pc \vdash_e \mathsf{delete}\, e : \dot{\tau} \qquad pc \sqcup \mathsf{lev}(\dot{\tau}) \sqsubseteq \dot{\tau}.\overline{f} = \sigma_f}{\Gamma, pc \vdash_e \mathsf{delete}\, e.f : \mathsf{PRIM}^{\sigma_f}}$$

$$\textsc{Function Literal}\quad \frac{\Gamma' = \mathsf{hoist}(\Gamma[x \mapsto \dot{\tau}.\mathsf{arg}, this \mapsto \dot{\tau}.\mathsf{this}], s) \qquad pc' = \dot{\tau}.\mathsf{wef} \quad \mathsf{lev}(\dot{\tau}) \sqcup pc \sqsubseteq pc' \quad \Gamma', pc', \dot{\tau} \vdash_s s \hookrightarrow s'}{\Gamma, pc \vdash_e \mathsf{function}\,(x)[\dot{\tau}]\{s\} \hookrightarrow \mathsf{function}\,(x)[\dot{\tau}]\{s'\} : \dot{\tau}}$$

$$\textsc{Static Method Call}\quad \frac{\forall_{i=1,2}\, \Gamma, pc \vdash_e e_i : \dot{\tau}_i \qquad \dot{\tau}_f = \pi(\dot{\tau}_1, f) \qquad \sigma = pc \sqcup \mathsf{lev}(\dot{\tau}_1) \sqcup \mathsf{lev}(\dot{\tau}_f) \qquad \sigma \sqsubseteq \dot{\tau}_f.\mathsf{wef} \qquad \dot{\tau}_1^{\sigma} \preceq \dot{\tau}_f.\mathsf{this} \qquad \dot{\tau}_2^{\sigma} \preceq \dot{\tau}_f.\mathsf{arg}}{\Gamma, pc \vdash_e e_1.f(e_2) : (\dot{\tau}_f.\mathsf{ret})^{\sigma}}$$

$$\textsc{Verified Expr Stmt}\quad \frac{\Gamma, pc \vdash_e e \hookrightarrow e' : \dot{\tau}}{\Gamma, pc, \dot{\tau}_{ret} \vdash_s e \hookrightarrow e'}$$

$$\textsc{Dyn. Expression Stmt}\quad \frac{e \in \mathcal{E}xpr_{\not\downarrow} \qquad s = @monitor(\Gamma, pc, \dot{\tau}_r, e)}{\Gamma, pc, \dot{\tau}_r \vdash_s e \hookrightarrow s}$$

$$\textsc{(Partially) Verified Conditional}\quad \frac{\Gamma, pc \vdash_e e \hookrightarrow e' : \dot{\tau} \qquad \forall_{i=0,1}\, \Gamma, pc \sqcup \mathsf{lev}(\dot{\tau}), \dot{\tau}_r \vdash_s s_i \hookrightarrow s'_i}{\Gamma, pc, \dot{\tau}_{ret} \vdash_s \mathsf{if}(e)\,\{s_1\}\,\mathsf{else}\,\{s_2\} \hookrightarrow \mathsf{if}(e')\,\{s'_1\}\,\mathsf{else}\,\{s'_2\}}$$

$$\textsc{Monitored Conditional}\quad \frac{e \in \mathcal{E}xpr_{\not\downarrow} \qquad s = @monitor(\Gamma, pc, \dot{\tau}_{ret}, \mathsf{if}(e)\,\{s_1\}\,\mathsf{else}\,\{s_2\})}{\Gamma, pc, \dot{\tau}_{ret} \vdash_s \mathsf{if}(e)\,\{s_1\}\,\mathsf{else}\,\{s_2\} \hookrightarrow s}$$

Fig. 2. Static typing core JS expressions

the explanations of Rules Literal, Variable, and Assignment as they are standard).

Static field projection. As a given field may be defined anywhere in the prototype chain of the inspected object, this rule needs to take into account the whole prototype chain of that object. To this end, we overload function π to model a static prototype chain inspection procedure. Informally, $\pi(\dot{\tau}, f)$ computes the *lub* between the security types of f in the prototype chain of objects of type $\dot{\tau}$ and upgrades the external level of this type with the *lub* between the existence levels of the field *_proto_* in that prototype chain.

Example 3 (Leaks via Prototype Mutations). The program below creates three empty objects bound to: `o0`, `o1`, and `o2`. Then, it creates a field named `public1` in both `o1` and `o2`, which is set to 0 in `o1` and to 1 in `o2`. Depending on the value of a *high* variable `secret`, the prototype of `o0` is either set to `o1` or to `o2`. Finally, the *low* variable `public1` is set to the value of the field `public1` of the prototype of `o0` (because `o0` does not define that field), thereby creating an implicit information flow between `secret` and `public`.

```
o0 = {}; o1 = {}; o2 = {}; o1.public1 = 0; o2.public1 = 1;
if(secret){o0._proto_ = o1} else {o0._proto_ = o2}; public = o0.public1
```

Letting Γ be the typing environment of Table 3, it follows that $\pi(\Gamma(\texttt{o0}), \texttt{public1}) = \textsc{prim}^H$ because $\Gamma(\texttt{o0}).\overline{\texttt{_proto_}} = H$. Hence, the assignment `public = o0.public1` is not typable as the type of `o0.public1`, $\textsf{PRIM}^H$, is not lower than or equal to $\textsf{PRIM}^L$.

Static Member Check. Since the domain of an object can change at execution time and since programs can check if a given field is defined using the keyword `in`, the mere existence of a field may disclose secret information. The existence security levels declared in object security types serve to control this type of information flows. However, analogously to field projections, this rule needs to take into account the whole prototype chain of the inspected object, because the field whose existence is being checked may be defined anywhere in that prototype chain. To this end, we make use of the static function $\bar{\pi}(\dot{\tau}, f)$ that computes the *lub* between the existence levels of f and *_proto_* in the prototype chain of objects of type $\dot{\tau}$.

Example 4 (Leaks via Membership Checks). The program below creates an object with two fields `secret1` and `secret2`. Then, depending on the value of a *high* variable `secret`, it deletes either `secret1` or `secret2` from the domain of `o`. Finally, the *low* variable `public` is assigned to `true` if `secret1` is defined in the prototype chain of `o` or to `false` if it is not, thereby creating an implicit flow between `secret` and `public`.

```
o = {}; o.secret1 = 0; o.secret2 = 0;
if (secret) { delete o.secret1 } else { delete o.secret2 }; public = secret1 in o
```

Letting Γ be the typing environment of Table 3, it follows that $\bar{\pi}(\Gamma(\texttt{o}), \texttt{secret1}) = \textsc{prim}^H$ because $\Gamma(\texttt{o}).\overline{\texttt{secret1}} = H$. Hence, the last assignment is not typable as the type of the expression `secret1 in o`, $\textsf{PRIM}^H$, is not lower than or equal to $\textsf{PRIM}^L$.

Static field assignment. The first constraint of the rule checks if the type of the assigned expression is a subtype of the assigned field type, thus preventing the assignment of a secret value to a public field. The second constraint checks if the context level is lower than or equal to the existence level of the assigned field, thereby preventing the creation of a visible field depending on secret information.

Field Deletion. The rule checks if the context level is lower than or equal to the field's existence level, thereby preventing visible fields from being deleted in invisible contexts.

Functional literal. This rule checks if the context level is lower than or equal to the writing effect of the type of the function literal, thereby preventing the evaluation of function literals that update or create *public* resources inside secret contexts. Then, the type system types the body of the function literal using the typing environment obtained by extending the current one with the type of the the formal argument, the type of the keyword this, and the types of the variables declared in the body of the function literal. To this end, we make use of a syntactic function hoist that extends the typing environment given as its first argument with the mappings from the variables declared in the statement given as its second argument to their respective security types. Note that this rule may re-write the the body of the function literal in order to enable the dynamic analysis.

Method call. This rule first verifies if the context level is lower than or equal to the writing effect of the method to call, thereby preventing the calling of a method that creates or updates *public* resources depending on *secret* values. Then, the rule checks if the type of the object on which the method is called and the type of the function argument match the type of the keyword this and the type of the formal parameter. The method call is finally typed with the return type of the method type upgraded with the context level.

Dyn. expression statement. This rule wraps every expression that contains a dynamic field operation inside a boundary statement. Recall that $\mathcal{E}xpr_{\lightning}$ denotes the set of Core JS dynamic field operations.

Conditional. If the conditional guard contains a dynamic field operation, the whole conditional is wrapped inside a boundary statement. In the opposite case, the type system types both branches, upgrading the context level with the external level of the security type of the conditional guard.

Example 5 (Hybrid versus Static Typing of the Running Example). Consider the program from Example 1 and the typing environment of Table 3. When typing the assignment public = o[g()], which contains a dynamic field operation, the type system applies the Dyn. expression statement rule and wraps the whole assignment inside a boundary statement. All the other statements, which do not contain dynamic field operations, are fully statically verified and, therefore, left unchanged. Hence, the resulting program is given by:

```
o = {}; o.secret1 = secret_input(); o.public1 = public_input();
o.public2 = public_input(); @monitor(@type_env, @pc, @ret, public = o[g()])
```

If, instead, the type system tried to statically type this assignment, it would need to check that the type of o[g()] was less than or equal to the type of public, PRIM^L. Since we do not know the value to which the call to g evaluates, the type system would need to use the *lub* between the types of all the fields declared in the type of o. Consequently, as one of those fields has type PRIM^H, the assignment would not be typable.

MONITOR SYNC
$$\frac{\Psi \xrightarrow{\alpha_l} \Psi' \quad \omega \xrightarrow{\alpha_l} \omega'}{(\!|\Psi, \Omega \cup \{\omega\}|\!) \rightarrow (\!|\Psi', \Omega \cup \{\omega'\}|\!)}$$

UNMONITORED STEP
$$\frac{\Psi \xrightarrow{\alpha_l} \Psi' \quad \forall_{\omega \in \Omega}\, \mathsf{er}(\omega) \neq l}{(\!|\Psi, \Omega|\!) \rightarrow (\!|\Psi', \Omega|\!)}$$

MONITOR CONFIGURATION +
$$\frac{l = \mathsf{head}(L) \qquad \forall_{\omega \in \Omega}\, \mathsf{er}(\omega) \neq l \qquad \omega' = \langle \Gamma, \dot{\tau}_r, l, pc :: [\,], [\,] \rangle}{(\!|\langle H, L, E[@monitor(\Gamma, \dot{\tau}_r, pc, s)] \rangle, \Omega|\!) \rightarrow (\!|\langle H, L, E[@monitor(s)] \rangle, \Omega \cup \{\omega'\}|\!)}$$

MONITOR CONFIGURATION - 1
$$\frac{\Psi = \langle H, L, E[@monitor(\underline{v})] \rangle \quad \mathsf{head}(L) = \mathsf{er}(\omega) \quad \Psi' = \langle H, L, E[\underline{v}] \rangle}{(\!|\Psi, \Omega \cup \{\omega\}|\!) \rightarrow (\!|\Psi', \Omega|\!)}$$

MONITOR CONFIGURATION - 2
$$\frac{\Psi = \langle H, L, E[@monitor(\mathsf{return}\, \underline{v})] \rangle \quad \mathsf{head}(L) = \mathsf{er}(\omega) \quad \Psi' = \langle H, L, E[\mathsf{return}\, \underline{v}] \rangle}{(\!|\Psi, \Omega \cup \{\omega\}|\!) \rightarrow (\!|\Psi', \Omega|\!)}$$

Fig. 3. Monitored semantics rules

4 Dynamic Typing Secure Information Flow in Core JS

The goal of a boundary statement is to enable and disable the information flow analysis at runtime. In this section, we define the semantics of the boundary operator by extending the semantics of Core JS with optional tracking of security types and verification of security constraints.

Monitored Semantics. A configuration of the monitored semantics has the form $(\!|\Psi, \Omega|\!)$ where Ψ is a Core JS configuration and Ω is a possibly empty set of monitor configurations. A monitor configuration ω is associated to a specific function call and has the form $\omega = \langle \Gamma, \dot{\tau}_r, l, o, \rho \rangle$ where: **(1)** Γ is a typing environment, **(2)** $\dot{\tau}_r$ is the type of the function that is executing, **(3)** l is the identifier of the environment record associated to the function call that is being monitored, **(4)** o is a *control context*, which is a list containing the levels of the expressions on which the monitored statement branched in order to reach the current context, and **(5)** ρ is an *expression context*, which is a list consisting of the security types of the values of the current evaluation context. The rules of the monitored semantics are given in Fig. 3 and described below. We use $\mathsf{er}(\omega)$ to denote the location of the environment record associated with ω.

Rule MONITOR SYNC corresponds to a monitored step. The transition of the monitor is synchronised with the transition of Core JS semantics through an *internal event* α_l, where l identifies the running function that performed a computation step.

Rule UNMONITORED STEP models the case where there is no matching monitor configuration for the current computation step. In this case, Core JS semantics performs an unconstrained computation step (that takes place outside a boundary statement).

Rule MONITOR CONFIGURATION + generates a new monitor configuration for verifying the statement inside a boundary statement. In order to account for computation steps inside boundary statements, we extend the syntax of evaluation contexts with a special boundary context: $E = @monitor(E')$.

LITERAL
$$\frac{\rho' = \mathsf{PRIM}^{\perp} :: \rho}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\mathsf{lit}_l}{\rightarrow} \langle o, \rho' \rangle}$$

THIS
$$\frac{\rho' = \Gamma(this) :: \rho}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\mathsf{this}_l}{\rightarrow} \langle o, \rho' \rangle}$$

VARIABLE
$$\frac{\rho' = \Gamma(x) :: \rho}{\Gamma, \dot{\tau}_r \vdash \langle o, \rho \rangle \overset{\mathsf{var}_l(x)}{\rightarrow} \langle o, \rho' \rangle}$$

VARIABLE ASSIGNMENT
$$\frac{pc = \mathsf{head}(o) \qquad \dot{\tau} = \mathsf{head}(\rho) \qquad \dot{\tau}^{pc} \preceq \Gamma(x)}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\mathsf{v\text{-}ass}_l(x)}{\rightarrow} \langle o, \rho \rangle}$$

FIELD PROJECTION
$$\frac{pc = \mathsf{head}(o) \qquad \rho = \dot{\tau}_2 :: \dot{\tau}_1 :: \rho' \qquad \dot{\tau} = \pi(\dot{\tau}_1, f) \qquad \sigma = pc \sqcup \mathsf{lev}(\dot{\tau}_1) \sqcup \mathsf{lev}(\dot{\tau}_2)}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\mathsf{f\text{-}proj}_l(f)}{\rightarrow} \langle pc :: o, \rho, \dot{\tau}^{\sigma} :: \rho' \rangle}$$

MEMBERSHIP CHECK
$$\frac{pc = \mathsf{head}(o) \qquad \rho = \dot{\tau}_2 :: \dot{\tau}_1 :: \rho' \qquad \sigma = \bar{\pi}(\dot{\tau}_1, f) \sqcup \mathsf{lev}(\dot{\tau}_1) \sqcup \mathsf{lev}(\dot{\tau}_2) \sqcup pc}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\mathsf{in}_l(f)}{\rightarrow} \langle o, \mathsf{PRIM}^{\sigma} :: \rho \rangle}$$

FIELD ASSIGNMENT
$$\frac{\rho = \dot{\tau}_3 :: \dot{\tau}_2 :: \dot{\tau}_1 :: \rho' \quad pc = \mathsf{head}(o) \qquad \sigma = \mathsf{lev}(\dot{\tau}_1) \sqcup \mathsf{lev}(\dot{\tau}_2) \sqcup pc \qquad \dot{\tau}_3^{\sigma} \preceq \dot{\tau}_1.f \quad \sigma \sqsubseteq \dot{\tau}_1.\overline{f}}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\mathsf{f\text{-}ass}_l(f)}{\rightarrow} \langle o, \dot{\tau}_3 :: \rho' \rangle}$$

FIELD DELETION
$$\frac{\rho = \dot{\tau}_2 :: \dot{\tau}_1 :: \rho' \qquad \sigma = \dot{\tau}_1.\overline{f} \qquad \mathsf{lev}(\dot{\tau}_1) \sqcup \mathsf{lev}(\dot{\tau}_2) \sqcup \mathsf{head}(o) \sqsubseteq \sigma}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\mathsf{del}_l(f)}{\rightarrow} \langle o, \mathsf{PRIM}^{\sigma} :: \rho' \rangle}$$

METHOD CALL
$$\frac{\rho = \dot{\tau}_3 :: \dot{\tau}_2 :: \dot{\tau}_1 :: \rho' \quad pc = \mathsf{head}(o) \qquad \dot{\tau}_f = \pi(\dot{\tau}_1, f) \quad \sigma = \mathsf{lev}(\dot{\tau}_1) \sqcup \mathsf{lev}(\dot{\tau}_2) \sqcup pc \qquad \sigma \sqsubseteq \dot{\tau}_f.\mathsf{wef} \quad \dot{\tau}_1^{\sigma} \preceq \dot{\tau}_f.\mathsf{this} \quad \dot{\tau}_3^{\sigma} \preceq \dot{\tau}_f.\mathsf{arg}}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\mathsf{m\text{-}call}_l(f)}{\rightarrow} \langle o, (\dot{\tau}_f.\mathsf{ret})^{\sigma} :: \rho \rangle}$$

IF - BRANCH
$$\frac{o' = \mathsf{lev}(\dot{\tau}) :: o}{\Gamma, \dot{\tau}_r, l \vdash \langle o, \dot{\tau} :: \rho \rangle \overset{\mathsf{if}_l}{\rightarrow} \langle o', \rho \rangle}$$

IF - END
$$\Gamma, \dot{\tau}_r, l \vdash \langle \sigma :: o, \rho \rangle \overset{\succ\mathsf{if}_l}{\rightarrow} \langle o, \rho \rangle$$

Fig. 4. Dynamic typing core JS expressions and statements

Rules MONITOR CONFIGURATION - 1 and MONITOR CONFIGURATION - 2 remove a monitor configuration from the current set of monitor configurations when its corresponding statement finishes executing.

Monitoring Rules. Monitor transitions are defined in Fig. 4. We use $\Gamma, \dot{\tau}_r, l \vdash \langle o, \rho \rangle \overset{\alpha_l}{\rightarrow} \langle o', \rho' \rangle$ as shorthand for $\langle \Gamma, \dot{\tau}_r, l, o, \rho \rangle \overset{\alpha_l}{\rightarrow} \langle \Gamma, \dot{\tau}_r, l, o', \rho' \rangle$. The constraints enforced by the monitor are the same as the constraints enforced by the type system of Sect. 3. However, in contrast to the type system, the monitor can precisely type dynamic expressions, since it has access to field names computed at runtime.

Example 6 (Monitoring a Dynamic Field Look-up). In the following, we present the sequence of monitor configurations generated when executing the statement: `@monitor(@type_env, @pc, @ret, public = o[g()])` (check the running example).

$$\langle \perp, [\,] \rangle \overset{\mathsf{var}_l(\mathtt{o})}{\rightarrow} \langle L, \dot{\tau}_o \rangle \overset{\mathsf{var}_l(\mathtt{g})}{\rightarrow} \langle L, \langle \dot{\tau}_{g._} \overset{H}{\rightarrow} \mathsf{PRIM}^L \rangle^L :: \dot{\tau}_o \rangle \overset{\mathsf{f\text{-}call}_l}{\rightarrow} \langle L, \mathsf{PRIM}^L :: \dot{\tau}_o \rangle$$

If `g()` returns `public1`: $\overset{\mathsf{f\text{-}proj}_l(\mathtt{public1})}{\rightarrow} \langle L, \mathsf{PRIM}^L \rangle \overset{\mathsf{v\text{-}ass}_l(\mathtt{public})}{\rightarrow} \langle L, \mathsf{PRIM}^L \rangle$

If `g()` returns `private1`: $\overset{\mathsf{f\text{-}proj}_l(\mathtt{private1})}{\rightarrow} \langle L, \mathsf{PRIM}^H \rangle \overset{\mathsf{v\text{-}ass}_l(\mathtt{public})}{\not\rightarrow}$

We consider two different cases: the case in which `g()` evaluates to `public1` and the case in which it evaluates to `secret1`. While in the first case, the execution is allowed to go through, in the second one it gets stuck, because the program tries to assign a secret value to a public variable.

Let us now briefly explain the rules that better illustrate our choices when designing the monitor. Since, by default, all literal values are public, when a *literal* value is evaluated, the monitor simply pushes $\mathsf{PRIM}^{\perp}$ onto the expression stack. In contrast, when a *variable* is evaluated, the monitor has to read its type from the typing environment and push it onto the expression stack. When a *field projection* is evaluated, the first two types on the expression stack are the types of the expressions that evaluate to the field name and to the inspected object, respectively. Furthermore, the name of the inspected field is available in the internal event that labels the transition. Hence, the monitor simply has to replace the first two types of the expression stack with the type of the inspected field upgraded with the external levels of the types of the current subexpressions. When an *if statement* is evaluated, the type of the conditional guard is on top of the expression stack. Hence, the monitor simply pops that type out of the expression stack and pushes its external level (upgraded with the current *pc*) onto the control stack. Complementarily, when the execution leaves the branch of a conditional, the monitor just pops out the top of the control stack.

Implementation. Instead of wrapping statements containing dynamic field operations within boundary statements, which are not part of the JavaScript language, the prototype of the hybrid type system [15] in-lines the monitoring logic in the statement itself [16]. This approach has the advantage of being immediately deployable. The prototype implementation was used to secure simple Web application accessible online [15].

5 Security Guarantees

This section describes the security guarantees offered by the proposed analysis. To formally define the absence of information leaks, we rely on an intuitive notion of *low-projection* [14] that establishes the part of a heap that an attacker at a given security level can see. Informally, given a heap H, an attacker at level σ can observe:

1. the existence of a field f in the domain of an object whose type has external level $\leq \sigma$ and associates f with an existence level $\leq \sigma$ and
2. the value of a field f in the domain of an object whose type has external level $\leq \sigma$ and associates f with a security type with external level $\leq \sigma$.

Figure 5 presents a labelled object together with its low-projection at level L. The object in the figure has three fields: f_1, f_2, and f_3. An attacker at level L can observe both the existence and the value of f_1 since it has *low* existence level and is associated with a visible value and the existence but not the value of f_2, since it has *low* existence level but is associated with an invisible value. The attacker

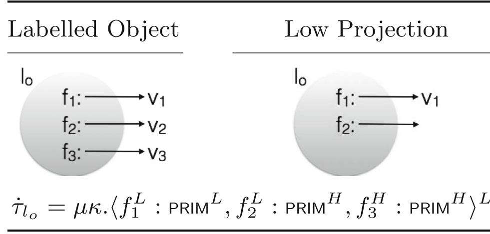

Fig. 5. A labelled object and its low projection

can neither observe the value nor the existence of f_3 because it has *high* existence level and is associated with an invisible value. Two heaps H_0 and H_1 are said to be *low-equal* at level σ, written $H_0 \sim_\sigma H_1$, if they coincide in their respective low-projections. Theorem 1 states that the monitored successfully-terminating execution of a program generated by the static type system on two low-equal heaps always yields two low-equal heaps. A sketch of the proof of Theorem 1 is given in the long-version of the paper available online at [15].

Theorem 1 (Noninterference). *For any typing environment Γ, levels σ and pc, security type $\dot{\tau}$, statement, s, and two heaps H_0 and H_1, such that $\Gamma, pc, \dot{\tau} \vdash_s s \hookrightarrow s'$, $H_0 \sim_\sigma H_1$, and $(\!|\langle H_i, [\,], s' \rangle, \{\}|\!) \rightarrow^* (\!|\langle H_i', [\,], v_i \rangle, \{\}|\!)$ for $i = 0, 1$, it holds that $H_0' \sim_\sigma H_1'$.*

6 Related Work

There is a wide variety of mechanisms for enforcing and verifying secure information flow, ranging from purely static type systems [14,18] to different flavours of dynamic analysis [2,13]. The main mechanisms for securing information flow in JavaScript [1,6,8] are mostly-dynamic due to the dynamicity of the language.

There is a long line of research on safety types for JavaScript which dates back to the seminal work of Thieman [17]. Since then, the TypeScript programming language [11] was proposed as a flexible language that adds optional types to JavaScript with the goal of harnessing the flexibility of real JavaScript, while at the same time providing some of the advantages otherwise reserved for statically typed languages, such as informative compiling errors. Recently, Rastogi et al. [12] designed and implemented a new gradual type system for safely compiling TypeScript to JavaScript. The soundness of the proposed approach is guaranteed by combining strict static checks with residual runtime checks. We believe that our work can serve as a basis for extending TypeScript types with security labels in order to verify secure information flow in TypeScript web applications.

Gradual type systems for secure information flow have been proposed for a pure lambda calculus [3] and for a core ML-like language with references [4]. The goal of these two works is significantly different from ours, as their main

intent is to cater for the use of *polymorphic* security labels. For instance, the type language proposed in [4] includes a special annotation "?" representing an unknown security level at static time. Expressions that use variables whose types contain the unknown level annotation, "?", cannot be precisely typed at static time. The programmer can introduce runtime casts in points where values of a pre-determined security type are expected. Then the dynamic analysis checks whether or not a cast can be *securely* performed during execution. However, in order to verify such casts at runtime, these analyses must track security labels during the execution of both dynamically verified and statically verified program regions. In contrast, our analysis only needs to dynamically verify the execution of program regions which were not statically verified.

7 Conclusions

We propose a sound hybrid typing analysis for enforcing secure information flow in a core of JavaScript that includes dynamic field operations. Furthermore, our analysis can be easily extended to handle other dynamic constructs of the language such as `eval` or unknown code, which only need to be wrapped inside the proposed boundary statement. Finally, we have implemented our analysis and used it to verify a web application available online [15].

This work follows a well-established trend on combining static and dynamic analysis to devise more permissive and efficient hybrid mechanisms [13]. Our approach can be applied to other scenarios, such as the verification of isolation properties [9], where it could be used to derive mostly-static lightweight enforcement mechanisms from prior purely static specifications.

Acknowledgments. We acknowledge funding from the EPSRC grant reference EP/K032089/1 (Fragoso Santos) and the ANR project AJACS ANR-14-CE28-0008 (Jensen, Rezk, and Schmitt). No new data was collected in the course of this research.

References

1. Bichhawat, A., Rajani, V., Garg, D., Hammer, C.: Information flow control in webkit's javascript bytecode. In: Abadi, M., Kremer, S. (eds.) POST 2014 (ETAPS 2014). LNCS, vol. 8414, pp. 159–178. Springer, Heidelberg (2014)
2. Devriese, D., Piessens, F.: Noninterference through secure multi-execution. In: SP (2010)
3. Disney, T., Flanagan, C.: Gradual information flow typing. In: STOP (2011)
4. Fennell, L., Thiemann, P.: Gradual security typing with references. In: CSF (2013)
5. Flanagan, C.: Hybrid type checking. In: POPL (2006)
6. De Groef, W., Devriese, D., Nikiforakis, N., Piessens, F.: Flowfox: a web browser with flexible and precise information flow control. In: CCS (2012)
7. Hedin, D., Birgisson, A., Bello, L., Sabelfeld, A.: JSFlow: tracking information flow in JavaScript and its APIs. In: SAC (2014)
8. Hedin, D., Sabelfeld, A.: Information-flow security for a core of JavaScript. In: CSF (2012)

9. Maffeis, S., Taly, A.: Language-based isolation of untrusted JavaScript. In: CSF (2009)
10. Matthews, J., Findler, R.B.: Operational semantics for multi-language programs. In: ACM TOPLAS (2009)
11. Microsoft. TypeScript language specification. Technical report, Microsoft (2014)
12. Rastogi, A., Swamy, N., Fournet, C., Bierman, G., Vekris, P.: Safe & efficient gradual typing for TypeScript. In: POPL (2015)
13. Russo, A., Sabelfeld, A.: Dynamic vs. static flow-sensitive security analysis. In: CSF (2010)
14. Sabelfeld, A., Myers, A.C.: Language-based information-flow security. IEEE J. Sel. Areas Commun. **21**(1), 5–19 (2003)
15. Santos, J.F.: Online materials - hybrid type system 2015. http://www.doc.ic.ac.uk/~jfaustin
16. Santos, J.F., Rezk, T.: An information flow monitor-inlining compiler for securing a core of javascript. In: Cuppens-Boulahia, N., Cuppens, F., Jajodia, S., Abou El Kalam, A., Sans, T. (eds.) SEC 2014. IFIP AICT, vol. 428, pp. 278–292. Springer, Heidelberg (2014)
17. Thiemann, P.: Towards a type system for analyzing javascript programs. In: Sagiv, M. (ed.) ESOP 2005. LNCS, vol. 3444, pp. 408–422. Springer, Heidelberg (2005)
18. Volpano, D.M., Irvine, C.E., Smith, G.: A sound type system for secure flow analysis. J. Comput. Secur. **4**(2), 167–187 (1996)
19. Wright, A., Felleisen, M.: A syntactic approach to type soundness. Inf. Comput. **115**(1), 38–94 (1994)