

Towards a Program Logic for JavaScript

Philippa Gardner

Imperial College London
pg@doc.ic.ac.uk

Sergio Maffei

Imperial College London
maffei@doc.ic.ac.uk

Gareth Smith

Imperial College London
gds@doc.ic.ac.uk

Abstract

JavaScript has become the most widely used language for client-side web programming. The dynamic nature of JavaScript makes understanding its code notoriously difficult, leading to buggy programs and a lack of adequate static-analysis tools. We believe that logical reasoning has much to offer JavaScript: a simple description of program behaviour, a clear understanding of module boundaries, and the ability to verify security contracts.

We introduce a program logic for reasoning about a broad subset of JavaScript, including challenging features such as prototype inheritance and `with`. We adapt ideas from separation logic to provide tractable reasoning about JavaScript code: reasoning about easy programs is easy; reasoning about hard programs is possible. We prove a strong soundness result. All libraries written in our subset and proved correct with respect to their specifications will be well-behaved, even when called by arbitrary JavaScript code.

Categories and Subject Descriptors F.3.1 [Theory of Computation]: Logics and Meanings of Programs—Specifying and Verifying and Reasoning about Programs

General Terms Reliability, Theory, Verification

Keywords JavaScript, Separation Logic, Web

1. Introduction

JavaScript has become the de-facto language for client-side web programming. Ajax web applications, used in e.g. Google Docs, are based on a combination of JavaScript and server-side programming. JavaScript has become an international standard called ECMAScript [14]. Adobe Flash, used in e.g. YouTube, features ActionScript, a programming language based on ECMAScript. Even web applications written in e.g. Java, F# or purpose-designed languages such as Flapjax or Milescript are either compiled to JavaScript, or they lack browser integration or cross-platform compatibility. JavaScript is currently the assembly language of the Web, and this seems unlikely to change.

JavaScript was initially created for small web-programming tasks, which benefited from the flexibility of the language and tight browser integration. Nowadays, the modern demands placed on JavaScript are huge. Although this flexibility and browser integration are still key advantages, the inherent dynamic nature of the language makes current web code based on ECMAScript 3 noto-

riously difficult to write and use [12, 18, 28]. The expectation is that ECMAScript 5 and future standards will improve the situation. However, although the main browsers now support ECMAScript 5, the majority of code being written today is in ECMAScript 3. Even if there is a wide acceptance of ECMAScript 5, which is certainly not clear from the current blogs, it is inevitable that ECMAScript 5 libraries will have to interface properly with ECMAScript 3 code. We therefore believe that there is a growing need for general-purpose, expressive analysis tools for both ECMAScript 3 and 5, which provide simple, correct descriptions of program behaviour and a clear understanding of module boundaries.

We introduce a program logic for reasoning about ECMAScript 3. While it is tempting to ignore the ‘ugly’ parts of the language, and reason only about ‘well-written’ code, in practice JavaScript programs have to interface with arbitrary web code. We therefore use a model of the language that does not shun the most challenging JavaScript features. For example, the behaviour of prototype inheritance, and the interplay between scoping rules and the `with` statement, is complex. This means that our basic reasoning rules must also be complex. We overcome this complexity by establishing natural layers of abstraction on top of our basic reasoning. With principled code, we can stay within these layers of abstraction and the reasoning is straightforward. With arbitrary code, we must break open the appropriate abstraction layers until we can re-establish the invariants of the abstraction. In this way, we are able to provide clean specifications of a wide variety of JavaScript programs.

Our reasoning is based on separation logic. Separation logic has proven to be invaluable for reasoning about programs which directly manipulate the heap, such as C and Java programs [2, 3, 8, 20, 36]. A key characteristic of JavaScript is that the entire state of the language resides in the object heap, in a structure that imperfectly emulates the variable store of many other programming languages. It is therefore natural to investigate the use of separation logic to verify JavaScript programs. In fact, we had to fundamentally adapt separation logic, both to present an accurate account of JavaScript’s emulated variable store (see Section 2: Motivating Examples) and also to establish soundness. For soundness, it is usual to require that all the program commands are ‘local’, according to a definition first given in [20]. Many JavaScript statements are not local by this definition: for example, even a simple variable read is non-local because its result may depend on the *absence* of certain fields in the emulated variable store. We instead prove soundness using a concept of *weak locality*, recently introduced by Smith [30].

In this paper, we reason about a substantial subset of JavaScript, including prototype inheritance, `with`, simple functions (no recursive or higher-order functions) and simple `eval`. Our treatment of functions and `eval` is precisely enough to expose fully the complexity of the emulated variable store. Building on the work of Charlton and Reus [6, 29], we will extend our reasoning to higher-order functions and complex `eval` in future. We prove soundness of our reasoning with respect to a faithful subset of the formal

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POPL’12, January 25–27, 2012, Philadelphia, PA, USA.
Copyright © 2012 ACM 978-1-4503-1083-3/12/01...\$10.00




Figure 1. A JavaScript emulated variable store

operational semantics of Maffeis *et al.* [18]. Our soundness result has powerful implications. Library code written in our subset and proved correct with respect to their specifications will be well-behaved, even when called by arbitrary JavaScript code. Our result is constructed in such a way that it will be simple to extend to reasoning about higher-order functions and complex `eval` in due course.

2. Motivating Examples

As convincingly argued in [12, 18, 27, 28], there are many reasons why the behaviour of JavaScript programs is complex. For a start, JavaScript is a dynamically typed, prototype-oriented language, which does not have a standard notion of variable store. Instead, JavaScript variables are best regarded as being stored in the heap, in a structure which imperfectly emulates the variable store. This structure consists of an abstract list of scope objects, called the *scope chain*, analogous to stack frames in other languages. Every scope object has a pointer to a linked list of prototypes, providing prototype-based inheritance. Since scope objects inherit data from their prototypes, the value of a variable cannot be resolved by a simple list traversal. Variable resolution is further complicated by the fact that JavaScript objects may share a common prototype.

JavaScript’s behaviour can make apparently simple programs deceptively counter-intuitive. Consider the code `C` defined below:

```
x = null; y = null; z = null;
f = function(w){x = v; v = 4; var v; y = v;};
v = 5; f(null); z = v;
```

What values should the variables `x`, `y` and `z` store at the end of the program? The correct answer is `undefined`, 4 and 5 respectively. We explain how this occurs as we walk through our reasoning.

In Section 6.2 we prove the following triple of this code:

$$\frac{\{ \text{store}_{\text{LS}}(x, y, z, f, v) * \mathbf{I} \doteq \text{LS} \}}{\text{C}} \left\{ \begin{array}{l} \exists L. \text{store}_{\text{LS}}(|x : \text{undefined}, y : 4, z : 5, f : L, v : 5) \\ * \mathbf{I} \doteq \text{LS} * \text{true} \end{array} \right\}$$

We distinguish a global logical expression \mathbf{I} with value LS denoting the scope chain. The store predicate $\text{store}_{\text{LS}}(x, y, z, f, v)$ states that the store-like structure referred to by LS contains *none* of the program variables mentioned; the variables occur to the left of the bar. The store predicate $\text{store}_{\text{LS}}(|x : \text{undefined}, y : 4, z : 5, f : L, v : 5)$ denotes the final values for all the variables; the variables occur to the right of the bar with assigned values.

To understand the complexity of the heap structures described by store predicates, consider the example heap given in Figure 1. This diagram illustrates a typical shape of a JavaScript variable store. Each object is denoted by a box. In this example, the current list of scope objects is given by $\mathbf{I} = [l_1, l_2, l_3, l_4, l_g]$, where the l_i are object addresses and l_g is a distinguished object con-

taining the global variables which must occur at the end of the current list of scope objects. Each scope object has a pointer to a list of prototypes, with the arrows representing prototype relationships. These prototype lists can be shared, but cannot form cycles. In ECMAScript 3, prototype lists must either end with the distinguished object l_{op} or be empty. However, many implementations (SpiderMonkey, V8 and WebKit) allow the programmer to directly access and change the prototype pointers, allowing incomplete prototype chains ending in `null` but not allowing the creation of cycles. We work with incomplete prototype chains, since we want ECMAScript 3 library code to work well with such implementations.

To look up the value of a variable `x` in our example heap, we check each object for a field with name `x`, starting with l_1 , checking the prototype list from l_1 then moving along the list of scope objects. In our example, the `x` in object l will be found first, since the whole prototype chain of l_2 will be visited before l_3 . When reading the value stored in `x`, this is all we need to know. If we write to the same variable `x`, the effect will be to create a new field `x` at l_2 . This new field will override the `x` field in object l in the usual prototype-oriented way.

All of this messy detail is abstracted away by the store predicate. The formation of this predicate is subtle and requires some adaptation of separation logic. As well as the separating conjunction $*$ for reasoning about disjoint heaps, we introduce the *sepish* connective \boxtimes for reasoning about partially separated heaps. It is used, for example, to account for the sharing of prototype lists illustrated in Figure 1. We also use the assertion $(l, x) \mapsto \emptyset$, which states that the field `x` is *not* present at object address l . This predicate is reminiscent of the ‘out’ predicate in [7] stating that values are not present in a concurrent list. It is necessary to identify the first `x` in the structure: in our example, the `x` at l is the first `x`, since it does not occur in the prototype list of l_1 nor in the prototype list of l_2 until l .

Our store predicate allows us to make simple inferences about variable assignments, without breaking our store abstraction:

$$\left\{ \begin{array}{l} \text{store}_{\text{LS}}(x, y, z, f, v) * \mathbf{I} \doteq \text{LS} \\ x = \text{null}; \\ \left\{ \begin{array}{l} \text{store}_{\text{LS}}(y, z, f, v | x : \text{null}) \\ * \mathbf{I} \doteq \text{LS} * \text{true} \end{array} \right\} \end{array} \right\}$$

where the assertion `true` hides potentially garbage-collected prototype lists.

The evaluation of the function expression `function(w) { ... }` has the effect of creating a new function object and returning the address L of that object. The object contains a number of internal housekeeping fields, including `@body` which contains the body of the function and `@scope` which stores the function closure LS . Our inference for the function definition is approximately:

$$\left\{ \begin{array}{l} \text{store}_{\text{LS}}(f, v | x : \text{null}, y : \text{null}, z : \text{null}) * \mathbf{I} \doteq \text{LS} \\ f = \text{function}(w) \{ \dots \} \\ \left\{ \begin{array}{l} \exists L. \text{store}_{\text{LS}}(v | x : \text{null}, y : \text{null}, z : \text{null}, f : L) * \\ (L, @body) \mapsto \lambda w. \{ \dots \} * (L, @scope) \mapsto \text{LS} * \\ \mathbf{I} \doteq \text{LS} * \text{true} \end{array} \right\} \end{array} \right\}$$

As well as the store predicate, we assert that the state also contains object cells such as $(L, @scope) \mapsto \text{LS}$. This assertion means that there is an object with address L in the heap, and it definitely contains at least the field `@scope` which has value LS . The assertion says nothing about any other field of L . We assert that our function object has fields `@body` and `@scope`. The full specification, given in Section 6.2, is actually a little more complicated than this. For now, we hide additional housekeeping fields in the assertion `true`.

We know that this program example is challenging, because the final values of the variables are counter-intuitive. All the complexity of the example occurs within the function call. When JavaScript calls a function, it performs two passes on the body: in the first pass, it creates a new scope object and initialises local variables to `undefined`; in the second pass, it runs the code in the newly constructed local scope. Our reasoning reflects this complexity. The Hoare triple for the function call has the following shape:

$$\left\{ \begin{array}{l} \text{store}_{\text{LS}}(|x : \text{null}, y : \text{null}, z : \text{null}, f : L, v : 5) * \\ (L, @body) \mapsto \lambda w. \{ \dots \} * (L, @scope) \mapsto \text{LS} * \\ \mathbf{I} \doteq \text{LS} * \text{true} \end{array} \right\} \quad \begin{array}{l} f(\text{null}); \\ \{ ??? \} \end{array}$$

To find a suitable postcondition, we must reason about the function body. The precondition of the function-body triple given below is constructed from the first pass of the function call. As well as containing the precondition of the function call, it contains a new scope object L' with fields given by the parameter of the function and the local variables discovered by the first pass. For our example, it contains the assertions $(L', w) \mapsto \text{null}$ for the parameter declaration and $(L', v) \mapsto \text{undefined}$ for the local variable declaration. The object L' also has a `@proto` field, which points to `null` since scope objects do not inherit any behaviour, and a `@this` field, which can only be read. We also have the predicate $\text{newobj}(L', @proto, @this, w, v)$, which asserts the absence from L' of all the fields we have not mentioned as parameters. Knowing this absence of fields is essential if, in the function body, we wish to write to variables, such as the `x` and `y`, which do not appear in the local scope object. Finally, the new scope object L' is prepended to the scope list \mathbf{l} .

Using this precondition, we are now able to give the triple obtained by the second pass of the function call:

$$\left\{ \begin{array}{l} \exists L'. \mathbf{l} \doteq L' : \text{LS} * \\ \text{store}_{\text{LS}}(|x : \text{null}, y : \text{null}, z : \text{null}, f : L, v : 5) * \\ (L, @body) \mapsto \lambda w. \{ \dots \} * (L, @scope) \mapsto \text{LS} * \\ \text{newobj}(L', @proto, @this, w, v) * \\ (L', w) \mapsto \text{null} * (L', v) \mapsto \text{undefined} * \\ (L', @proto) \mapsto \text{null} * (L', @this) \mapsto L'' * \text{true} \\ x = v ; v = 4 ; \text{var } v ; y = v ; \end{array} \right\} \quad \left\{ \begin{array}{l} \exists L'. \mathbf{l} \doteq L' : \text{LS} * \\ \text{store}_{\text{LS}}(|x : \text{undefined}, y : 4, z : \text{null}, f : L, v : 5) * \\ (L, @body) \mapsto \lambda w. \{ \dots \} * (L, @scope) \mapsto \text{LS} * \\ \text{newobj}(L', @proto, @this, w, v) * \\ (L', w) \mapsto \text{null} * (L', v) \mapsto 4 * \\ (L', @proto) \mapsto \text{null} * (L', @this) \mapsto L'' * \text{true} \end{array} \right\}$$

The postcondition follows simply from the three assignments in the function body: first, variable `x` gets the value `undefined`, since this is the current value of the local `v`; then the local `v` is assigned `4`; and, finally, the global variable `y` is assigned the value of the local variable `v`. The `var v` statement has no effect in the second pass of the function call.

The postcondition of the function call is the postcondition of the function body, with local scope object L' popped off the current scope list \mathbf{l} :

$$\left\{ \begin{array}{l} \exists L'. \text{store}_{\text{LS}}(|x : \text{undefined}, y : 4, z : \text{null}, f : L, v : 5) * \\ (L, @body) \mapsto \lambda w. \{ \dots \} * (L, @scope) \mapsto \text{LS} * \\ \text{newobj}(L', @proto, @this, w, v) * \\ (L', w) \mapsto \text{null} * (L', v) \mapsto 4 * (L', @proto) \mapsto \text{null} * \\ (L', @this) \mapsto L'' * \mathbf{l} \doteq \text{LS} * \text{true} \end{array} \right\}$$

Reasoning about the final assignment is straightforward, with `z` assigned the value of the global variable `v`. The final postcondition is

obtained using the consequence rule to hide the function object and local scope object behind the assertion `true`, since they are surplus to requirements, and existentially quantifying function object L :

$$\left\{ \begin{array}{l} \exists L. \text{store}_{\text{LS}}(|x : \text{undefined}, y : 4, z : 5, f : L, v : 5) * \\ \mathbf{l} \doteq \text{LS} * \text{true} \end{array} \right\}$$

Part of the challenge of understanding this example is knowing the scope of local variable `v`. In JavaScript, variables can only be declared local to functions, not other blocks such as `if` and `while`. This can lead to undesirable behaviour, especially when a local variable overloads the name of a global variable. One controversial technique for solving this problem is to use the `with` statement and a literal object to declare local variable blocks precisely where they are needed. Using `with` is often considered bad practice, and it is deprecated in ECMAScript 5. However, it is widely used in practice [22, 28] and can certainly be used to improve the program readability. We are able to reason about even extremely confusing uses of `with`. Consider the program C' :

```
a = {b:1}; with (a){f=function(c){return b}};
a = {b:2}; f(null)
```

Armed with an operational understanding of JavaScript's emulated variable store, it is not so difficult to understand that this program returns the value `1`, even though the value of `a.b` at the end of the program is `2`. It may not be quite so clear that this program can fault. It may also execute arbitrary code from elsewhere in the emulated variable store, leading to a possible security violation.

We only understood this example properly by doing the verification. In Section 6.2, we prove the triple:

$$\left\{ \text{store}_{\mathbf{l}}(a, f) \right\} \text{true} \mapsto \emptyset \text{true} \quad \left\{ \begin{array}{l} (l_{op}, f) \mapsto \emptyset \text{true} \\ (l_{op}, @proto) \mapsto \text{null} \end{array} \right\} \quad \left\{ \mathbf{r} \doteq \mathbf{l} * \text{true} \right\}$$

A similar proof is possible for a precondition where `a` and `f` are in the store with arbitrary values. Either precondition ensures the program returns the value `1` as expected. The obvious first try was to have e.g. just $\text{store}_{\mathbf{l}}(a, f)$ as the precondition. This does not work as, when reasoning about the assignment to the variable `f`, we cannot assert that the variable `f` is not in the local scope. As discussed earlier, we work with incomplete prototype chains in the emulated variable store. With some web code, it is possible for programmers to directly access and change the prototype pointers, resulting in the distinguished object l_{op} not being a part of the emulated variable store. This means that l_{op} may contain the field `f` without violating our proposed precondition. The statement `a = {b:1}` then results in the creation of a new object L , with no field `f` as expected, but with field `@proto` pointing to l_{op} which may contain `f`:

$$\exists L. (L, b) \mapsto 1 * (L, f) \mapsto \emptyset * (L, @proto) \mapsto l_{op}$$

The `with` statement makes this new object L the most local cell in our emulated variable store. But because `f` is an inherited property of L , the meaning of the assignment to `f` has now changed – it has become an overriding assignment of a local variable. We write the new function `f` into the most local object L . When the program returns from the `with` statement, the function call to `f(null)` will fault, since there is no `f` in scope to call. In many web browsers (including Chrome, Firefox, and Safari), this behaviour can be induced by running the program `Object.prototype.f = 4 ; window.__proto__ = null ; C'` in the default starting state¹. The `__proto__` notation allows the programmer to directly access and change the internal `@proto`

¹ In web browsers, `Object.prototype` usually contains a pointer to the object l_{op} and `window` usually contains a pointer to the object l_g .

fields, resulting in a program state which causes the program C' to fault. Non-standard features such as this are widely used in practice, and so it is important that our reasoning be robust to the states they produce. In this example, it is possible to induce a similarly tricky starting state without the use of any non-standard features, which results in the call to $f()$ executing arbitrary code not mentioned in our program. Consider running the program `Object.prototype.f = function(c) {C''};C''` where C'' is suspicious code. This could result in a security breach if f is passed sensitive data.

3. Operational Semantics

We define a big-step operational semantics for a large subset of JavaScript that represents faithfully the inheritance, prototyping and scoping mechanisms described in the ECMAScript 3 standard. Our semantics follows closely the full small-step JavaScript semantics of Maffei, Mitchell and Taly [18], except that we make some simplifications as discussed in Section 3.5. We work with a big-step semantics because it connects better to our reasoning.

3.1 Heaps

The JavaScript heap is a partial function $H: \mathcal{R} \rightarrow \mathcal{V}$ that maps *references*, $r \in \mathcal{R} = \mathcal{L} \times \mathcal{X}$, which are pairs of memory locations and field names, to values. This structure emphasises the important role that references play in the semantics of the language. Values $v \in \mathcal{V}$ can be basic values \mathcal{V} , locations l and lambda abstractions $\lambda x. e$. The set of locations \mathcal{L} is lifted to a set $\mathcal{L}_{\text{null}}$ containing the special location `null`, analogous to a null-pointer in C, which cannot be in the domain of any heap. We denote the empty heap by emp , a heap cell by $(l, x) \mapsto v$, the union of two disjoint heaps by $H_1 * H_2$, and a read operation by $H(l, x)$.

An object is represented by a set of heap cells addressed by the same location but with different field names. For ease of notation, we use $l \mapsto \{x_1: v_1, \dots, x_n: v_n\}$ as a shorthand for the object $(l, x_1) \mapsto v_1 * \dots * (l, x_n) \mapsto v_n$.

As discussed, JavaScript has no variable store. Instead, variables are resolved with respect to a scope object implicitly known at run time. Scope objects are just objects whose locations are recorded in the *scope chain* (we use a standard notation $[], e:L, L \# L$ for lists). Each scope object has a pointer to a *prototype list* (which need not point to l_{op}). A variable x is resolved as the property named “ x ” of the first object in the scope chain whose prototype list defines “ x ”. Scoping constructs, such as function calls and `with`, cause sub-expressions to be evaluated with respect to a local scope object, by putting the local scope object at the beginning of the scope chain and then removing it after the sub-expressions have been evaluated. All user programs are evaluated starting from the default scope chain $[l_g]$, where l_g is the location of the global JavaScript object, described below. The final object in any scope chain is always l_g , but duplicates in a scope chain are allowed. For example, the JavaScript program `with(window){C}` is perfectly valid, and (usually) results in the subprogram C being evaluated in a state in which l_g is both the most global *and* the most local object.

The auxiliary scope function σ , defined below, returns the location of the first object in the scope chain to define a given variable. It depends on the prototype function π , which returns the location of the first object in the prototype chain to define the variable.

Scope and prototype resolution: $\sigma(H, l, x)$ and $\pi(H, l, x)$.

$$\sigma(H, [], x) \triangleq \text{null}$$

$$\frac{\pi(H, l, x) \neq \text{null}}{\sigma(H, l:L, x) \triangleq l} \quad \frac{\pi(H, l, x) = \text{null}}{\sigma(H, l:L, x) \triangleq \sigma(H, L, x)}$$

$$\pi(H, \text{null}, x) \triangleq \text{null}$$

$$\frac{(l, x) \in \text{dom}(H)}{\pi(H, l, x) \triangleq l} \quad \frac{(l, x) \notin \text{dom}(H) \quad H(l, @proto) = l'}{\pi(H, l, x) \triangleq \pi(H, l', x)}$$

The set of variable names \mathcal{X} is partitioned in two disjoint sets: the *internal* variable names \mathcal{X}^I and the *user* variable names \mathcal{X}^U . The internal names $\mathcal{X}^I \triangleq \{@scope, @body, @proto, @this\}$ are not directly accessible by user-defined ECMAScript 3 code, but are used by the semantics. As noted in Section 2, some implementations do give programmers direct access to some of these internal variables, but we do not model those implementation-dependant programs. We do model all the program states that such programs might create, and take care not to make strong assumptions about the types or values of those suspect internal variables. In this way, libraries verified using our techniques are robust in the presence of unverified real-world code. User names are denoted by $x, y, z \in \mathcal{X}^U$ and are considered a subset of strings; keywords such as `var` are not valid variable names. It is worth anticipating at this point a subtlety of the JavaScript semantics. The evaluation of a user variable x does not return its value, but rather the reference $l \cdot x$ where such value can be found (l' is obtained using the σ predicate). In general, the values $r \in \mathcal{V}^R$ returned by JavaScript expressions can be normal values \mathcal{V} or references \mathcal{R} . When a user variable x needs to be dereferenced in an expression, the semantics implicitly calls the dereferencing function γ , defined below, which returns the value denoted by the reference.

Dereferencing values: $\gamma(H, r)$.

$$\frac{r \neq l \cdot x}{\gamma(H, r) \triangleq r} \quad \frac{\pi(H, l, x) = \text{null} \quad l \neq \text{null}}{\gamma(H, l \cdot x) \triangleq \text{undefined}} \quad \frac{\pi(H, l, x) = l' \quad l' \neq \text{null}}{\gamma(H, l \cdot x) \triangleq H(l', x)}$$

We introduce the definition of a well-formed JavaScript heap. A JavaScript expression can only be evaluated in a sub-heap of a well-formed heap, with respect to a valid scope chain. All JavaScript expressions maintain the well-formedness of heaps. A heap H is *well-formed* (denoted by $H \vdash \diamond$) if its objects and prototype chains are well-formed (prototype chains must be acyclic, they can end with `null` rather than l_{op}), and if it contains the global scope object l_g and the distinguished objects l_{op} and l_{fp} , which we will see later serve as the default prototypes for new objects and functions. A scope chain L is *valid* with respect to heap H , denoted by $\text{schain}(H, L)$, if all the locations in the chain correspond to objects allocated in H , and if it ends with the global object l_g . Formal definitions are given in [11].

The default initial state H^0 is the smallest well-formed heap that also contains the special function object l_e (the `eval` function) and its prototype l_{ep} :

$$H^0 \triangleq \left(\begin{array}{l} \text{obj}(l_g, l_{op}) * \text{obj}(l_{op}, \text{null}) \\ * \text{obj}(l_{fp}, l_{op}) * \text{obj}(l_e, l_{ep}) * \text{obj}(l_{ep}, l_{op}) \end{array} \right)$$

where $\text{obj}(l, l')$ denotes an object at location l which is empty except for the prototype l' :

$$\text{obj}(l, l') \triangleq (l, @proto) \mapsto l'$$

We conclude this section by defining the heap update $-[-]$ operation which will be used by the semantics. Many JavaScript expressions update a heap cell if it is present, or create it if it is not present. We reflect this form of update in our heap update operation.

Update $H[(l, x) \mapsto v]$.

$$\frac{(l, x) \notin \text{dom}(H) \quad l \neq \text{null}}{H[(l, x) \mapsto v] \triangleq H * (l, x) \mapsto v} \quad \frac{}{(H * (l, x) \mapsto v)[(l, x) \mapsto v'] \triangleq H * (l, x) \mapsto v'} \quad \frac{}{H[(\text{null}, x) \mapsto v] \triangleq H[(l_g, x) \mapsto v]}$$

The last rule says that an update to a non-allocated reference (whose location is `null`) has the effect of allocating and updat-

ing a new global variable, which mimics JavaScript's behaviour of implicitly creating global variables when a variable name is first used.

3.2 Values and Expressions

To aid our presentation, we introduce the syntax of all JavaScript programs, statements and expressions simply as expressions. Our operational semantics will only describe the behaviour of well-formed programs, as defined by [18].

Syntax of terms: values v and expressions e .

v	$::=$	$n \mid m \mid \text{undefined} \mid \text{null}$
e	$::=$	$e; e \mid x \mid v \mid \text{if}(e)\{e\}\{e\} \mid \text{while}(e)\{e\} \mid \text{var } x$ $\mid \text{this} \mid \text{delete } e \mid e \oplus e \mid e.x \mid e(e) \mid e = e$ $\mid \text{function}(x)\{e\} \mid \text{function } x(x)\{e\} \mid \text{new } e(e)$ $\mid \{x_1 : e_1 \dots x_n : e_n\} \mid e[e] \mid \text{with}(e)\{e\}$

where $\oplus \in \{+, -, *, /, \&\&, ||, ==, .\}$

A basic value v can be a number n , a string m (including the user variable names), the special constant `undefined` or the null location. The operator \oplus denotes a standard number and boolean operator, or string concatenation. Expressions e include sequential composition, variable lookup, literal values, conditional expressions, loops, local variable declaration, `this`, deletion, arithmetic and string concatenation, object property lookup, function call and eval, assignment, function declaration, recursive functions, constructors, literal objects, computed access and the `with` statement.

3.3 Evaluation Rules

An expression e is evaluated in a heap H , with respect to a scope chain L . If it successfully terminates, it returns a modified heap H' and a final value r . Selected evaluation rules are given below and then discussed. See [11] for the full definition. Recall that the set of variables is a subset of the set of strings, that a heap value v can be a basic value v , a memory location l or a function closure $\lambda x.e$, and that a return value r can also be a reference $l.x$.

Operational semantics: $H, L, e \longrightarrow H', r$.

Notation: $H, L, e \xrightarrow{\gamma} H', v \triangleq \exists r. (H, L, e \longrightarrow H', r \wedge \gamma(H', r) = v)$.

(Definition)	$\frac{H, L, e \longrightarrow H', v}{H, L, \text{var } e \longrightarrow H', \text{undefined}}$	(Value)	$\frac{H, L, v \longrightarrow H, v}{H, L, v \longrightarrow H, v}$
(Member Access)	$\frac{H, L, e \xrightarrow{\gamma} H', l \quad l \neq \text{null}}{H, L, e.x \longrightarrow H', l.x}$	(Computed Access)	$\frac{H, L, e1 \xrightarrow{\gamma} H1, l \quad l \neq \text{null}}{H1, L, e2 \xrightarrow{\gamma} H', x} \quad \frac{H1, L, e2 \xrightarrow{\gamma} H', x}{H, L, e1[e2] \longrightarrow H', l.x}$
(Variable)	$\frac{\sigma(H, L, x) = l}{H, L, x \longrightarrow H, l.x}$	(Object)	$\frac{H_0 = H * \text{obj}(l, l_{op}) \quad \forall i \in 1..n. \left(\begin{array}{l} H_{i-1}, L, e1 \xrightarrow{\gamma} H'_i, v_i \\ H_i = H'_i[(l, x_i) \mapsto v_i] \end{array} \right)}{H, L, \{x_1 : e1, \dots, x_n : e_n\} \longrightarrow H_n, l}$
(Binary Operators)	$\frac{H, L, e1 \xrightarrow{\gamma} H'', v1 \quad H'', L, e2 \xrightarrow{\gamma} H', v2 \quad v1 \oplus v2 = v}{H, L, e1 \oplus e2 \longrightarrow H', v}$	(Assignment)	$\frac{H, L, e1 \longrightarrow H1, l.x \quad H1, L, e2 \xrightarrow{\gamma} H2, v \quad H' = H2[(l, x) \mapsto v]}{H, L, e1=e2 \longrightarrow H', v}$
(This)	$\frac{\sigma(H, L, @this) = l_1 \quad \pi(H, l_1, @this) = l_2 \quad H(l_2, @this) = l'}{H, L, \text{this} \longrightarrow H, l'}$	(Function)	$\frac{H' = H * \text{obj}(l, l_{op}) * \text{fun}(l', L, x, e, l)}{H, L, \text{function}(x)\{e\} \longrightarrow H', l'}$

(Function Call)

$$\frac{H, L, e1 \longrightarrow H1, r1 \quad \text{This}(H1, r1) = l_2 \quad \gamma(H1, r1) = l_1 \quad l_1 \neq l_e \quad H1(l_1, @body) = \lambda x.e3 \quad H1(l_1, @scope) = L' \quad H1, L, e2 \xrightarrow{\gamma} H2, v \quad H3 = H2 * \text{act}(l, x, v, e3, l_2) \quad H3, l:L', e3 \xrightarrow{\gamma} H', v'}{H, L, e1(e2) \longrightarrow H', v'}$$

(Eval)

$$\frac{H, L, e1 \xrightarrow{\gamma} H1, l_e \quad H1, L, e2 \xrightarrow{\gamma} H2, m \quad \text{parse}(m) = e \quad H2, L, e \xrightarrow{\gamma} H', v'}{H, L, e1(e2) \longrightarrow H', v'}$$

(With)

$$\frac{H, L, e1 \xrightarrow{\gamma} H1, l \quad H1, l:L, e2 \longrightarrow H', r \quad \sigma, \pi, \gamma, \text{obj defined earlier} \quad \text{fun, this, act defined below}}{H, L, \text{with}(e1)\{e2\} \longrightarrow H', r}$$

We briefly discuss some of the evaluation rules that show non-standard features typical of JavaScript. Rule (Definition) for `var` simply executes e and throws away the return value. The `var` declaration is only used by `defs` (defined below) to identify function local variables. Rule (Variable) uses σ to determine the scope object where a given variable can be found, without de-referencing the variable. Rules (Member/Computed Access) return a reference to the object field denoted by the corresponding expressions. Rule (Object) uses the `obj` notation introduced in Section 3.1 to introduce a fresh, empty object at location l , and then initialises its fields accordingly. Freshness is ensured by well-formedness of H and disjointness of $*$. Rule (Binary Operators) assumes the existence of a semantic version \oplus for each syntactic operator \oplus . Each \oplus is a partial function, defined only on arguments of a basic type (in this case numbers or strings) and returning results of some other (possibly the same) basic type, corresponding to the intended meaning of the operation. Rule (Assignment) is quite subtle. Suppose we have the expression `x=4`. Now consider Figure 1 of Section 2. If x were defined as a field of the object l_1 then `x=4` would be an *overwriting* assignment. The value of the field x in the object l_1 would be overwritten with the value 4. If x were not found anywhere, then it would be created as a global variable – a field of the object l_g . Finally, consider the actual case in this figure: x is found to be a field of object l , which is a prototype of l_2 , which is in the scope chain. In this case, `x=4` is an *overriding* assignment, with the effect of creating a new field x in the object l_2 to override the existing x in l . This complexity is handled in two places. Firstly, the variable rule uses the σ function (defined in Section 3.1) to return a reference $l_2.x$. Note that the σ function does not return a reference to $l.x$ precisely because we wish to model this behaviour. Secondly, the heap update operation $H_2[\dots]$ manages the business of either overwriting an existing field or, in this case, creating a new field. Rule (This) resolves the `this` identifier. As we will see, when executing a method of an object, we use the internal variable `@this` to store the location of that object. When executing a function call, `@this` points to the global object l_g . The (This) rule uses σ and π to retrieve the value of `@this` which is then returned by the `this` statement. Rule (Function) introduces the notation $\text{fun}(l', L, x, e, l) \triangleq$

$$l' \mapsto \{\text{@proto}: l_{fp}, \text{prototype}: l, \text{@scope}: L, \text{@body}: \lambda x.e\}$$

to allocate a fresh function object at location l' . The internal prototype of the new function object points to the standard function prototype object l_{fp} . The rule also creates a new empty object at l and stores a pointer to it in the `prototype` field of the new function. If the function is ever used as a constructor (using the `new` keyword) then the object pointed to by this `prototype` field will serve as the prototype of the newly constructed object. Note that the field is mutable, and so may be used by a programmer to emulate a class-like inheritance structure.

Recall from Section 2 that JavaScript function calls can be surprisingly complex. We now describe rule (Function Call), which uses two auxiliary functions `This` and `act`. Recall also that the rule (`This`) uses internal `@this` fields to determine the semantics of the `this` keyword. The values of the `@this` fields are determined by the (Function Call) rule using the auxiliary function `This`:

$$\begin{aligned} \text{This}(H, l \cdot x) &\triangleq l \quad [(l, @this) \notin \text{dom}(H)] \\ \text{This}(H, r) &\triangleq l_g \quad [\text{otherwise}] \end{aligned}$$

To understand `This`, first notice that every newly created local scope object has a `@this` field, while no other objects ever will². The (Function Call) rule finds a pointer to a function in the location $l \cdot x$. If l is a regular object (which has no `@this` field), then the function must be a method of that object, and so the `@this` field of our new local scope object should point to l . On the other hand, if l is a special local scope object (which has a `@this` field), then the function must be a regular function (and not a method), and so the `@this` of our new local scope object should point to l_g . This unique behaviour precisely captures the behaviour of the ECMAScript 3 `this` keyword. The auxiliary function `act` describes the allocation of a new local scope object: $\text{act}(l, x, v, e, l') \triangleq$

$$l \mapsto \{x: v, @this: l', @proto: \text{null}\} * \text{defs}(x, l, e)$$

The object is allocated at address l , and contains a function parameter x with value v , the internal fields `@this`, `@proto` and the local variables declared in an expression e . The auxiliary function `defs`, defined in [11], searches the function body e for instances of the `var` keyword, and sets all the appropriate fields of our new local scope object to the value `undefined`, as discussed in Section 2. Rule (Eval) looks like an ordinary function call, but the function being called is not an ordinary function object. It is the special built-in object l_e . It assumes a partial function `parse` that parses a string m into a JavaScript expression e , only if there are no syntax errors.

The control expressions are mostly standard (see [11]), except for the unusual (With) rule that evaluates e_2 in a scope-chain starting with the object obtained by evaluating e_1 .

3.4 Safety

An important sanity property of the evaluation relation is that it preserves well-formedness of the heap, for any valid scope chain.

Theorem 1 (Well-Formedness). *Let H, L be such that $H \vdash \diamond$ and $\text{schain}(H, L)$. If $H, L, e \longrightarrow H', r$ then $H' \vdash \diamond$.*

Although the theorem is about end-to-end well-formedness, its proof (reported in [11]) shows that starting from a well-formed state and scope chain, all the intermediate states and scope chains visited during the computation are also well-formed, and all the locations occurring in intermediate return values correspond to objects effectively allocated on the heap.

3.5 JavaScript Subset

We work with a subset of JavaScript, in order to limit the size and complexity of our semantics for this paper. Our subset is substantial and, despite some minor omissions and simplifying assumptions discussed below, faithful to the ECMAScript 3 standard. A significant property of our semantics is that our programs will run reliably in states generated by any valid JavaScript program (including those reached by programs using non-standard features that we do not model, such as `__proto__`), or getters and setters. Our reasoning of Section 5 will therefore interface well with real-world JavaScript programs.

We do not model implicit type-coercion functions. Adding them is straightforward but would add significantly to the length of our

² l_g has `@this` in some, not all, implementations; [14] is silent on the issue.

presentation. We have no boolean type. Instead, where control structures (`if` and `while`) require a boolean, we use other types whose behaviour is equivalent to the type conversion that occurs in JavaScript. Values such as 0 and `null` behave like `false` and values such as 1 and `"string"` behave like `true`. For simplicity, we use an implicit return statement for functions. Moreover, our functions take only one parameter, rather than the arbitrary list of parameters usual in JavaScript, and do not have the `arguments` object or the `constructor` property. As mentioned in Section 3.2, we simplify our presentation of JavaScript programs, statements and expressions, into a single class of expressions. We also omit several JavaScript constructs such as labels, `switch` and `for`, as they do not contribute significantly to the problem of program reasoning. In this presentation we only consider the core language ECMAScript 3, and do not model the many standard libraries which browsers make available to programmers. Instead of exceptions, we have a single error condition denoted `fault`. Our reasoning conservatively avoids faults. This means that programs which are proved using our fault-avoiding local Hoare reasoning will run without throwing exceptions in JavaScript interpreters.

4. Assertion Language

Our assertion language follows that of Parkinson and Bierman [4, 23, 24], in their work on reasoning about Java. They use assertions of the form $(l, x) \mapsto v$ to denote that a partial heap contains the object l with field x which has value v . Using the separating conjunction $*$ [20], the assertion $((l, x) \mapsto v) * ((l, y) \mapsto w)$ declares that a heap contains an object l which must have two separate fields x, y with the appropriate values. The assertion $(l, x) \mapsto v * (l, x) \mapsto w$ is unsatisfiable since it declares two fields x for l .

This style of reasoning is not enough for JavaScript. We must also assert negative information about fields not being in the heap, and extend the structural assertions of separation logic to account for partial separation due to shared prototype chains. Recall the example of a JavaScript emulated variable store in Figure 1. To find the value of x in the store, we must not only determine that the object l has a field x but also determine that no fields named x occur earlier in the emulated store. We use assertions $(l, x) \mapsto \emptyset$ to declare that a heap contains an object l which *does not* have a field x . The assertion $(l, x) \mapsto \emptyset * (l, y) \mapsto w$ declares that the heap contains an object l which does not have field x but does have field y . The assertions $(l, x) \mapsto v * (l, x) \mapsto \emptyset$ and $(l, x) \mapsto \emptyset * (l, x) \mapsto \emptyset$ are unsatisfiable. Thus, the assertion $(l, x) \mapsto \emptyset$ states the full knowledge that field x is not in object l .

Now consider what happens when we want to describe the state of more than one variable at a time. In Section 5 we introduce a predicate σ which allows us to assert, for example, “The variable x is found in the store in object l_2 ”: $\sigma(-, l, x, l_2)$ or “The variable y is not in the store”: $\sigma(-, l, y, \text{null})$. Both of these assertions must make use of the `@proto` fields in order to navigate the variable store, so we cannot separate them using $*$. But the first assertion does not mention any y field, and the second assertion does not mention any x field, so we cannot join them with \wedge . In order to make both assertions about the same variable store, we need a way for them to share their common parts. To do this, we introduce the *sepish connective* $P \boxtimes Q$ which allows *partial* separation between heaps. We can use \boxtimes to describe the state of more than one variable in a given store: $\sigma(-, l, x, l_2) \boxtimes \sigma(-, l, y, \text{null})$, and we shall see in Section 5 that it is also invaluable when defining the σ predicate.

An assertion may be satisfied by a triple (h, L, ϵ) consisting of an abstract heap h , a scope chain L , and a logical environment ϵ . An *abstract JavaScript heap* is a partial function $H: \mathcal{R} \rightarrow \mathcal{V} \cup \{\emptyset\}$ that maps references, $r \in \mathcal{R} = \mathcal{L} \times \mathcal{X}$ to values $v \in \mathcal{V}$ or \emptyset . Abstract heaps thus declare information about fields not being present in an object, as well as the fields that are present. We also define an

evaluation function $[-]$ which takes an abstract heap to a concrete heap:

$$[h](l, x) \triangleq h(l, x) \text{ iff } (l, x) \in \text{dom}(h) \wedge h(l, x) \neq \emptyset$$

We use this function in Section 5 to define the relationship between our reasoning using Hoare triples and our operational semantics.

We define a logical environment ϵ , which is a partial function from logical variables $X \in \mathcal{X}^L$ to logical values \mathcal{V}^L , which may be a return value r , any expression e , \emptyset or a list Ls of logical values. We also define logical expressions E , which are different from program expressions in that they can not read or alter the heap. Expressions E are evaluated in a logical environment ϵ with respect to a current scope chain L .

Logical expressions and evaluation: $\llbracket E \rrbracket_\epsilon^L$.

$v \in \mathcal{V}^L ::= e \mid v \mid \emptyset \mid Ls$	$\epsilon : \mathcal{X}^L \rightarrow \mathcal{V}^L$
$E ::=$	
X	Logical variables
\mathbf{l}	Scope list
v	Logical values
$E \oplus E$	Binary Operators
$E : E$	List cons
$E \cdot E$	Reference construction
$\lambda E. E$	Lambda values

$$\llbracket v \rrbracket_\epsilon^L \triangleq v \quad \llbracket \mathbf{l} \rrbracket_\epsilon^L \triangleq L \quad \llbracket X \rrbracket_\epsilon^L \triangleq \epsilon(X)$$

$$\llbracket E_1 : E_2 \rrbracket_\epsilon^L \triangleq \llbracket E_1 \rrbracket_\epsilon^L : Ls \text{ if } \llbracket E_2 \rrbracket_\epsilon^L = Ls$$

$$\llbracket E_1 \cdot E_2 \rrbracket_\epsilon^L \triangleq l \cdot x \text{ if } \llbracket E_1 \rrbracket_\epsilon^L = l \wedge \llbracket E_2 \rrbracket_\epsilon^L = x$$

$$\llbracket E_1 \oplus E_2 \rrbracket_\epsilon^L \triangleq v \oplus v' \text{ if } \llbracket E_1 \rrbracket_\epsilon^L = v \wedge \llbracket E_2 \rrbracket_\epsilon^L = v'$$

$$\llbracket \lambda E_1. E_2 \rrbracket_\epsilon^L \triangleq \lambda x. \llbracket E_2 \rrbracket_\epsilon^L \text{ if } \llbracket E_1 \rrbracket_\epsilon^L = x$$

Assertions include the boolean assertions, structural assertions of separation logic and our new sepish connective, basic assertions for describing cells in a JavaScript heap, expression equality, set and list assertions, and quantification over logical variables.

Assertions.

$P ::=$		
$P \wedge P \mid P \vee P \mid \neg P \mid \text{true} \mid \text{false}$		Boolean assertions
$P * P \mid P \multimap P \mid P \boxtimes P$		Structural assertions
$(E, E) \mapsto E \mid \emptyset$		JavaScript heap assertions
$E = E$		Expression equality
$E \in \text{SET}$		Set membership
$E \in E$		List element
$\exists X. P \mid \forall X. P$		Quantification

Notation: $E \not\in E \triangleq \neg(E \in E)$ for $\in \in \{=, \in\}$

$$E_1 \circ E_2 \triangleq E_1 \circ E_2 \wedge \emptyset \text{ for } \circ \in \{=, \neq, \in, \notin\}.$$

The logical operators bind in order: $\neg, \exists, \forall, *, \wedge, \vee, \multimap$.

The structural assertions $*$ and \multimap are standard separation logic assertions. The separating conjunction $P * Q$ says that the heap may be split into two disjoint heaps, one satisfying P and the other Q . The right adjoint $P \multimap Q$ says that, whenever the heap is extended by a heap satisfying P , then the resulting heap satisfies Q . It is useful in proving the soundness of some of our layers of abstraction in Section 6, as given in [11]. The sepish connective $P \boxtimes Q$ is novel. It says that the heap may be split into two heaps, one satisfying P and the other Q , but these two heaps need not be disjoint. They may share zero or more common cells. We shall see in Section 5 that this is particularly useful when reasoning about the emulated variable store. It is possible to define \boxtimes analogously with \multimap , but since this is not useful for JavaScript reasoning we omit it here. Note that $P \wedge Q \Rightarrow P \boxtimes Q$ and $P * Q \Rightarrow P \boxtimes Q$, but neither of the reverse implications hold. The assertion $(E_1, E_2) \mapsto E_3$ declares information about a cell, including the information that field E_2 does not occur in object E_1 . Assertion \emptyset says that the heap is empty. The notation SET denotes a literal set, or a named set such as \mathcal{X} , the set of JavaScript field names. Note that since

there are sets for numbers, strings and locations, we can use set inclusion to assert the type of a particular JavaScript value.

The satisfaction relation $h, L, \epsilon \models P$ is defined below. The cases for the boolean assertions are not given as they are standard.

Satisfaction of assertions: $h, L, \epsilon \models P$.

$h, L, \epsilon \models P * Q$	$\iff \exists h_1, h_2. h \equiv h_1 * h_2 \wedge (h_1, L, \epsilon \models P) \wedge (h_2, L, \epsilon \models Q)$
$h, L, \epsilon \models P \multimap Q$	$\iff \forall h_1. (h_1, L, \epsilon \models P) \wedge h \# h_1 \implies ((h * h_1), L, \epsilon \models Q)$
$h, L, \epsilon \models P \boxtimes Q$	$\iff \exists h_1, h_2, h_3. h \equiv h_1 * h_2 * h_3 \wedge (h_1 * h_3, L, \epsilon \models P) \wedge (h_2 * h_3, L, \epsilon \models Q)$
$h, L, \epsilon \models (E_1, E_2) \mapsto E_3$	$\iff h \equiv (\llbracket E_1 \rrbracket_\epsilon^L, \llbracket E_2 \rrbracket_\epsilon^L) \mapsto \llbracket E_3 \rrbracket_\epsilon^L$
$h, L, \epsilon \models \emptyset$	$\iff h = \text{emp}$
$h, L, \epsilon \models E_1 = E_2$	$\iff \llbracket E_1 \rrbracket_\epsilon^L = \llbracket E_2 \rrbracket_\epsilon^L$
$h, L, \epsilon \models E \in \text{SET}$	$\iff \llbracket E \rrbracket_\epsilon^L \in \text{SET}$
$h, L, \epsilon \models E_1 \in E_2$	$\iff \llbracket E_1 \rrbracket_\epsilon^L \text{ is in the list } \llbracket E_2 \rrbracket_\epsilon^L$
$h, L, \epsilon \models \exists X. P$	$\iff \exists v. h, L, [\epsilon X \leftarrow v] \models P$
$h, L, \epsilon \models \forall X. P$	$\iff \forall v. h, L, [\epsilon X \leftarrow v] \models P$

We have given a direct definition of the sepish connective. When logical variables range over heaps, it can be derived: $P \boxtimes Q \iff \exists R. (R * P) * (R \multimap Q) * R$. It remains to be seen what natural logical properties are satisfied by this connective.

5. Program Reasoning

In the spirit of separation logic, we give small axioms and inference rules which precisely capture the behaviour of all JavaScript expressions, except for the usual approximation for `while` and conservative approximations of function call and `eval`. Reasoning about function call and `eval` is interesting and complex, and will be the focus of future work as outlined in Section 8. Because our reasoning captures the full complexity of JavaScript semantics, particularly with respect to the emulated variable store and the `with` command, it is possible to prove properties about extremely subtle programs. Unfortunately, proving any program at this level of abstraction involves a level of detail which most programmers would understandably rather avoid. We therefore provide several *layers of abstraction* which make it possible to reason at a natural high level about well-behaved JavaScript programs. We discuss these abstractions in Section 6.

Our fault-avoiding Hoare triples take the form: $\{P\}e\{Q\}$, which means “if e is executed in a state satisfying P , then it will not fault, and if it terminates it will do so in a state satisfying Q ”. The postcondition Q may refer to the special variable \mathbf{r} , which is equal to the return value of e .

5.1 Auxiliary Predicates

For our reasoning rules, we require predicates that correspond to functions, such as σ , π and γ , used by the operational semantics in Section 3.3. These predicates present us with two distinct challenges: sharing and precision. To consider sharing, recall the scope function σ from Section 3.1 when searching for a variable y in the example emulated store given in Figure 1. Since y is not present in the emulated store, σ will check the entire structure before returning `null` to indicate that y cannot be found. What is of interest to us is the order in which the cells in the store will be checked. Notice that the cell l will be checked twice, and the cell l_{op} will be checked three times. This is because the cells l and l_{op} are shared between the prototype chain of the object l_2 and the object l_4 . In addition, l_{op} is shared by these prototype chains and the prototype chain of l_9 . As we shall see below, we can describe these partially shared structures using our sepish connective \boxtimes . To consider precision, recall the function σ when this time searching for a variable x either in the emulated store of Figure 1, or in the identical store but for the omission of a single object such that the prototype of l_2 would

be l . In each of these stores, the σ function will return the same value – the location l_2 . In our program reasoning, our σ predicate (which will correspond to the σ function of the operational semantics) must be more precise as we need to distinguish between these two possible cases. With the challenges of sharing and precision in mind, we first give, then explain, the logical predicates σ , π and γ .

Logical predicates: σ, π, γ .

$$\begin{aligned} \sigma([], [], \text{null}) &\triangleq \bigcirc \\ \sigma([Ls], L' : Ls', X, L') &\triangleq \exists L. \pi(Ls, L', X, L) * L \neq \text{null} \\ \sigma((Ls_1 : Ls_2), L' : Ls', X, L) &\triangleq \\ &\pi(Ls_1, L', X, \text{null}) \boxtimes \sigma(Ls_2, Ls', X, L) \\ \pi([], \text{null}, -, \text{null}) &\triangleq \bigcirc \\ \pi([L], L, X, L) &\triangleq \exists V. (L, X) \mapsto V * V \neq \emptyset \\ \pi((L' : Ls), L', X, L) &\triangleq \\ &\exists N. (L', X) \mapsto \emptyset * (L', @proto) \mapsto N * \pi(Ls, N, X, L) \\ \gamma([], Val, Val) &\triangleq Val \notin \mathcal{R} \\ \gamma(Ls, L \cdot X, \text{undefined}) &\triangleq \pi(Ls, L, X, \text{null}) * L \neq \text{null} \\ \gamma(Ls, L_1 \cdot X, Val) &\triangleq \\ &\exists L_2. \pi(Ls, L_1, X, L_2) \boxtimes (L_2, X) \mapsto Val * Val \neq \emptyset \end{aligned}$$

These predicates closely follow the structure of the functions defined in 3.1, using \boxtimes to manage the challenge of sharing prototype chains mentioned earlier. The predicate $\sigma(-, L, x, l)$ holds only for abstract heaps h such that the function from Section 3.1 gives us $\sigma([h], L, x) = l$, meaning that the value of the variable x in the emulated variable store given by the list L can be found in the object (or a prototype of the object) at address l . However, recall the challenge of precision mentioned earlier. In order to distinguish between all the possible heaps which satisfy the predicate $\sigma(-, L, x, l)$, we work with an additional first argument. The predicate $\sigma(Ls, L, x, l)$ is “precise” in the sense that, for any abstract heap h , it holds for at most one subheap of h . The first argument Ls is a list of lists, which specifies the exact cells which must be visited (and the order in which they must be checked) in order to determine which object in the emulated variable store defines the variable x . For example, recall the heap illustrated in Figure 1. If the prototype of the cell l_2 has address l' , then the predicate $\sigma([l_1], [l_2, l', l], \mathbf{l}, x, l_2)$ is satisfied by the abstract heap consisting of the x and $@proto$ fields of the four objects in the lower left corner of that diagram. Notice that we do not need to visit every object in the variable store in order to discover the location of the variable x . In the spirit of the small axioms of separation logic, our predicate holds for the smallest possible heap in which we can be sure of discovering the variable we are interested in.

The predicates π and γ are similar in that they mirror their operational counterparts, with the addition of one extra argument to make them precise. In the case of π and γ , the first argument Ls is simply a list of addresses, rather than a list of lists of addresses, because each predicate only walks down at most one prototype chain.

The inference rules also require logical predicates corresponding to a number of other auxiliary functions given in Section 3. Below, we define `newobj` and `fun` predicates, which assert the existence of a fresh object and function object, and `decls` that returns the local variables of an expression. In order to reason about function call, we also use a `defs` predicate which we define in [11].

Auxiliary predicates

$$\begin{aligned} \text{This}(L \cdot -, L) &\triangleq (L, @this) \mapsto \emptyset \quad \text{where } L \neq l_g \\ \text{This}(L \cdot -, l_g) &\triangleq \exists V. (L, @this) \mapsto V * V \neq \emptyset \\ \text{True}(E) &\triangleq E \notin \{0, \text{“”}, \text{null}, \text{undefined}\} \\ \text{False}(E) &\triangleq E \in \{0, \text{“”}, \text{null}, \text{undefined}\} \\ \text{newobj}(L, V_1, \dots, V_n) &\triangleq \textcircled{*} V \in \mathcal{X} \setminus \{V_1 \dots V_n\} (L, V) \mapsto \emptyset \end{aligned}$$

$$\begin{aligned} \text{fun}(F, Env, X, Body, Proto) &\triangleq \\ (F, @scope) &\mapsto Env * (F, @body) \mapsto \lambda X. Body * \\ (F, \text{prototype}) &\mapsto Proto * (F, @proto) \mapsto l_{fp} \end{aligned}$$

$$\text{decls}(X, L, e) \triangleq \mathbf{x}_1, \dots, \mathbf{x}_n \quad \text{where } (L, \mathbf{x}_i) \in \text{dom}(\text{defs}(X, L, e))$$

5.2 Inference Rules

We define below some inference rules $\{P\}e\{Q\}$ for reasoning about JavaScript expressions. The full list can be found in [11].

Inference rules: $\{P\}e\{Q\}$.

$$\begin{aligned} \text{(Definition)} & \frac{\{P\}e\{Q\} \quad \mathbf{r} \notin \text{fv}(Q)}{\{P\}\text{var } e\{Q * \mathbf{r} \doteq \text{undefined}\}} & \text{(Value)} & \frac{}{\{\bigcirc\}\mathbf{v}\{\mathbf{r} \doteq \mathbf{v}\}} \\ \text{(Variable)} & \frac{P = \sigma(Ls_1, \mathbf{l}, \mathbf{x}, L) \boxtimes \gamma(Ls_2, L \cdot \mathbf{x}, V)}{\{P\}\mathbf{x}\{P * \mathbf{r} \doteq L \cdot \mathbf{x}\}} & \text{(Variable Null)} & \frac{P = \sigma(Ls, \mathbf{l}, \mathbf{x}, \text{null})}{\{P\}\mathbf{x}\{P * \mathbf{r} \doteq \text{null} \cdot \mathbf{x}\}} \\ \text{(Member Access)} & \frac{\{P\}e\{Q * \mathbf{r} \doteq V\} \quad Q = R * \gamma(Ls, V, L) * L \neq \text{null}}{\{P\}e.\mathbf{x}\{Q * \mathbf{r} \doteq L \cdot \mathbf{x}\}} \\ \text{(Computed Access)} & \frac{\{P\}e_1\{R * \mathbf{r} \doteq V_1\} \quad R = S_1 * \gamma(Ls_1, V_1, L) * L \neq \text{null} \quad \{R\}e_2\{Q * X \in \mathcal{X}^U * \mathbf{r} \doteq V_2\} \quad Q = S_2 * \gamma(Ls_2, V_2, X)}{\{P\}e_1[e_2]\{Q * \mathbf{r} \doteq L \cdot X\}} \\ \text{(Object)} & \frac{\forall i \in 1..n. (P_i = R_i * \gamma(Ls_i, Y_i, X_i) \quad \{P_{i-1}\}e_i\{P_i * \mathbf{r} \doteq Y_i\}) \quad Q = \left(P_n * \exists L. \left(\begin{array}{l} \text{newobj}(L, @proto, \mathbf{x}_1, \dots, \mathbf{x}_n) * \\ (L, \mathbf{x}_1) \mapsto X_1 * \dots * (L, \mathbf{x}_n) \mapsto X_n * \\ (L, @proto) \mapsto l_{op} * \mathbf{r} \doteq L \end{array} \right) \right)}{\mathbf{x}_1 \neq \dots \neq \mathbf{x}_n \quad \mathbf{r} \notin \text{fv}(P_n) \quad \{P_0\}\{\mathbf{x}_1 : e_1, \dots, \mathbf{x}_n : e_n\}\{Q\}} \\ \text{(Binary Operators)} & \frac{\{P\}e_1\{R * \mathbf{r} \doteq V_1\} \quad R = S_1 * \gamma(Ls_1, V_1, V_3) \quad \{R\}e_2\{Q * \mathbf{r} \doteq V_2\} \quad Q = S_2 * \gamma(Ls_2, V_2, V_4) \quad V = V_3 \oplus V_4}{\{P\}e_1 \oplus e_2\{Q * \mathbf{r} \doteq V\}} \\ \text{(Assign Global)} & \frac{\{P\}e_1\{R * \mathbf{r} \doteq \text{null} \cdot X\} \quad \{R\}e_2\{Q * (l_g, X) \mapsto \emptyset * \mathbf{r} \doteq V_1\} \quad Q = S * \gamma(Ls, V_1, V_2)}{\{P\}e_1 = e_2\{Q * (l_g, X) \mapsto V_2 * \mathbf{r} \doteq V_2\}} \\ \text{(Assign Local)} & \frac{\{P\}e_1\{R * \mathbf{r} \doteq L \cdot X\} \quad \{R\}e_2\{Q * (L, X) \mapsto V_3 * \mathbf{r} \doteq V_1\} \quad Q = S * \gamma(Ls, V_1, V_2)}{\{P\}e_1 = e_2\{Q * (L, X) \mapsto V_2 * \mathbf{r} \doteq V_2\}} \\ \text{(Function)} & \frac{Q = \left(\begin{array}{l} \exists L_1, L_2. \text{newobj}(L_1, @proto) * (L_1, @proto) \mapsto l_{op} * \\ \text{newobj}(L_2, @proto, \text{prototype}, @scope, @body) * \\ \text{fun}(L_2, \mathbf{l}, \mathbf{x}, e, L_1) * \mathbf{r} \doteq L_2 \end{array} \right)}{\{\bigcirc\}\text{function } (\mathbf{x}) \{e\}\{Q\}} \\ \text{(While)} & \frac{\{P\}e_1\{S * \mathbf{r} \doteq V_1\} \quad S = R * \gamma(Ls, V_1, V_2) \quad \{S * \text{True}(V_2)\}e_2\{P\} \quad Q = S * \text{False}(V_2) * \mathbf{r} \doteq \text{undefined} \quad \mathbf{r} \notin \text{fv}(R)}{\{P\}\text{while}(e_1)\{e_2\}\{Q\}} \\ \text{(With)} & \frac{\{P * \mathbf{l} \doteq L\}e_1\{S * \mathbf{l} \doteq L * \mathbf{r} \doteq V_1\} \quad S = R * \gamma(Ls, V_1, L_1) \quad \{S * \mathbf{l} \doteq L_1 : L\}e_2\{Q * \mathbf{l} \doteq L_1 : L\} \quad \mathbf{l} \notin \text{fv}(P, Q, R)}{\{P * \mathbf{l} \doteq L\}\text{with}(e_1)\{e_2\}\{Q * \mathbf{l} \doteq L\}} \end{aligned}$$

(Function Call)	
$\{P\}e1\{R_1 * \mathbf{r} \doteq F_1\}$	
$R_1 = \left(\begin{array}{l} S_1 \boxtimes \text{This}(F_1, T) \boxtimes \gamma(Ls_1, F_1, F_2)* \\ (F_2, @body) \mapsto \lambda X. e3 * (F_2, @scope) \mapsto Ls_2 \end{array} \right)$	
$\{R_1\}e2\{R_2 * \mathbf{l} \doteq Ls_3 * \mathbf{r} \doteq V_1\} \quad R_2 = S_2 * \gamma(Ls_4, V_1, V_2)$	
$R_3 = \left(\begin{array}{l} R_2 * \exists L. \mathbf{l} \doteq L : Ls_2 * (L, X) \mapsto V_2* \\ (L, @this) \mapsto T * \\ (L, @proto) \mapsto \text{null} * \text{defs}(X, L, e3)* \\ \text{newobj}(L, @proto, @this, X, \text{decls}(X, L, e3)) \end{array} \right)$	
$\{R_3\}e3\{\exists L. Q * \mathbf{l} \doteq L : Ls_2\} \quad \mathbf{l} \notin \text{fv}(Q, R_2)$	
$\{P\}e1(e2)\{\exists L. Q * \mathbf{l} \doteq Ls_3\}$	
(Frame)	(Consequence)
$\frac{\{P\}e\{Q\}}{\{P * R\}e\{Q * R\}}$	$\frac{\{P_1\}e\{Q_1\} \quad P \implies P_1 \quad Q_1 \implies Q}{\{P\}e\{Q\}}$
(Elimination)	(Disjunction)
$\frac{\{P\}e\{Q\}}{\{\exists X. P\}e\{\exists X. Q\}}$	$\frac{\{P_1\}e\{Q_1\} \quad \{P_2\}e\{Q_2\}}{\{P_1 \vee P_2\}e\{Q_1 \vee Q_2\}}$

Although most of the rules correspond closely to their operational counterparts, some rules deserve further comment. Rule (Definition) shows the use of the reserved variable \mathbf{r} to record the result of an expression. Rule (Variable) shows the use of \boxtimes to express the overlapping footprints of predicates σ and γ . Rule (Assign Global) shows the use of \emptyset to assert that certain known memory cells are available for allocation. Rule (Function Call) describes JavaScript's dynamic functions but does not support higher order reasoning. Rule (Frame) does not have the usual side condition because JavaScript stores all its program variables on the heap, so any variable modified by an expression is necessarily contained entirely within the footprint of the expression. Rules (Consequence), (Elimination) and (Disjunction) are standard.

5.3 Soundness

We show that our inference rules are sound with respect to the operational semantics of Section 3.3. When proving the soundness of any system involving the frame rule, it is usual to first show the locality of the programming language, and use that property to show the soundness of the frame rule [20]. Unfortunately, many JavaScript statements are not local according to this standard definition. We therefore use the recently introduced notion of weak locality from Smith's thesis [30].

Definition 2 (Soundness of a Hoare triple). A Hoare triple $\{P\}e\{Q\}$ is sound if, for all abstract heaps h , scope chains L and environments ϵ , it satisfies the following two properties:

$$\text{Fault Avoidance} : h, L, (\epsilon \setminus \mathbf{r}) \models P \implies [h], L, \mathbf{e} \not\rightarrow \text{fault}$$

$$\text{Safety} : \forall H, r. h, L, (\epsilon \setminus \mathbf{r}) \models P \wedge [h], L, \mathbf{e} \longrightarrow H, r \implies \exists h'. H = [h'] \wedge h', L, [\epsilon | \mathbf{r} \leftarrow r] \models Q.$$

Theorem 3 (Soundness). *All derivable Hoare triples $\{P\}e\{Q\}$ are sound according to Definition 2.*

The proof (reported in [11]) involves showing that the predicates used by the Hoare rules correspond to the auxiliary functions used by the semantics, showing that all JavaScript expressions are weakly local with respect to their preconditions, and finally showing that all our inference rules are sound.

6. Layers of Abstraction

As mentioned in Section 5, using the rules given so far, reasoning about JavaScript program involves detail that most programmers need never consider. Most of the time, programmers will work at a higher level of abstraction, for example, treating the emulated variable store as if it were a regular variable store. This is

a good practice as long as the abstraction holds. However, if the program happens to come across a corner case that breaks the abstraction, its resulting behaviour can appear almost inexplicable. This may be a particular problem when writing library code, since the programmer has no control over the programming discipline of the client who uses the library. In this section, we introduce several explicit abstraction formalisms. Many alternatives are possible, but those presented here are enough to demonstrate the concept and reason about some interesting programs. We are able to use these formalisms to reason at a comfortable, high-level way about many well-behaved programs. Crucially, we know exactly what the boundaries of these layers of abstraction are, so we can ensure that our programs remain safely within the abstraction. If we wish, we can even choose to temporarily break an abstraction, execute some particularly subtle low-level code, re-establish the abstraction and then continue to work at the high level.

6.1 Layer 1: Exploring the Scope List

Central to reasoning about JavaScript variables are the σ and π predicates. The first abstraction layer consists of alternative versions of these predicates which make reasoning about certain common cases simpler. The $\bar{\sigma}$ predicate unrolls from the end (the most global end) of the scope rather than from the beginning (the local end) which makes modifying a variable easier to specify. It makes use of \emptyset which says that a variable does not exist in a particular partial scope. The ϕ^{lg} predicate does the same, but excludes l_g from its specification, in order to make reasoning about global variable instantiation simpler. We give several useful triples about variable assignment. In [11], we prove these triples by using $\bar{\sigma}$ and showing the equivalence of σ and $\bar{\sigma}$.

Layer 1 Predicates.

$$\begin{aligned} \bar{\sigma}(Ls, Ls', Var, \text{null}) &\triangleq \emptyset(Ls, Ls', Var, \text{null}) \\ \bar{\sigma}(Ls_1 \# (Ls_2 : []), Ls', Var, L) &\triangleq \\ &\emptyset(Ls_1, Ls', Var, L) * \exists L_2. \pi(Ls_2, L, Var, L_2) * L_2 \neq \text{null} \\ \emptyset([], L' : Ls', -, L') &\triangleq \emptyset \\ \emptyset(Ls_2 : Ls, L' : Ls', Var, End) &\triangleq \\ &\pi(Ls_2, L', Var, \text{null}) \boxtimes \emptyset(Ls, Ls', Var, End) \\ \phi^{lg}([], L' : Ls', -, L') &\triangleq \emptyset \\ \phi^{lg}(Ls_2 : Ls, L' : Ls', Var, End) &\triangleq \\ &\pi^{lg}(Ls_2, L', Var, \text{null}) \boxtimes \phi^{lg}(Ls, Ls', Var, End) \\ \pi^{lg}([], \text{null}, -, \text{null}) &\triangleq \emptyset \\ \pi^{lg}([], l_g, -, \text{null}) &\triangleq \emptyset \\ \pi^{lg}([L'], L', Var, L') &\triangleq \exists V. (L', Var) \mapsto V * V \neq \emptyset * L' \neq l_g \\ \pi^{lg}((L' : Ls), L', Var, L) &\triangleq \exists N. (L', Var) \mapsto \emptyset * \\ &(L', @proto) \mapsto N * \pi^{lg}(Ls, N, Var, L) * L' \neq l_g \end{aligned}$$

These predicates give us much more flexibility to reason about JavaScript variables found in various places in the emulated variable store. Even at this quite low level, it is possible to prove general specifications about programs with many corner cases. A good example of this sort of reasoning is simple assignment statements. We prove the following general triples about simple assignments. The first three triples deal with the assignment of a constant to a variable, in the cases of variable initialisation, variable override, and variable overwrite respectively. The fourth triple deals with assigning the value of one variable to another. All four are proved sound in [11].

Simple assignments.

$$\begin{aligned} P &= \sigma(L_1 \# ((l_g : L_2) : L_3), \mathbf{l}, \mathbf{x}, \text{null}) \\ Q &= \left(\begin{array}{l} \exists L'_1, L'_3, Ls', G. \emptyset^{lg}(L'_1, \mathbf{l}, \mathbf{x}, l_g) \boxtimes \\ \pi(L_2, G, \mathbf{x}, \text{null}) \boxtimes \emptyset^{lg}(L'_3, Ls', \mathbf{x}, \text{null}) * (l_g, \mathbf{x}) \mapsto v * \\ (l_g, @proto) \mapsto G * \mathbf{l} \doteq _ \# (l_g : Ls') * \mathbf{r} \doteq v \end{array} \right) \\ \{P\}x &= v\{Q\} \end{aligned}$$

$$\begin{array}{l}
P = \sigma(L_1 \# [L:L_2], \mathbf{l}, \mathbf{x}, L) \text{ \# } (L, \mathbf{x}) \mapsto \emptyset \text{ \# } \gamma(L:L_2, L \cdot \mathbf{x}, V) \\
Q = \left(\frac{\exists L'. \phi(L_1, \mathbf{l}, \mathbf{x}, L) * (L, \mathbf{x}) \mapsto \mathbf{v} * (L, @proto) \mapsto Pr *}{\pi(L_2, Pr, \mathbf{x}, L') \text{ \# } (L', \mathbf{x}) \mapsto V * \mathbf{r} \doteq \mathbf{v}} \right) \\
\{P\}\mathbf{x} = \mathbf{v}\{Q\} \\
P = \sigma(L_1 \# [[L]], \mathbf{l}, \mathbf{x}, L) \text{ \# } (L, \mathbf{x}) \mapsto V * V \neq \emptyset \\
Q = \phi(L_1, \mathbf{l}, \mathbf{x}, L) * (L, \mathbf{x}) \mapsto \mathbf{v} * \mathbf{r} \doteq \mathbf{v} \\
\{P\}\mathbf{x} = \mathbf{v}\{Q\} \\
P = \left(\frac{\sigma(L_{s1}, \mathbf{l}, \mathbf{y}, L_y) \text{ \# } \gamma(L_{s2}, L_y \cdot \mathbf{y}, V_y) \text{ \# } \sigma(L_1 \# [[L]], \mathbf{l}, \mathbf{x}, L) \text{ \# } (L, \mathbf{x}) \mapsto V * V \neq \emptyset}{\sigma(L_{s1}, \mathbf{l}, \mathbf{y}, L_y) \text{ \# } \phi(L_1, \mathbf{l}, \mathbf{x}, L) * (L, \mathbf{x}) \mapsto V_y * \mathbf{r} \doteq V_y} \right) \\
Q = \sigma(L_{s1}, \mathbf{l}, \mathbf{y}, L_y) \text{ \# } \phi(L_1, \mathbf{l}, \mathbf{x}, L) * (L, \mathbf{x}) \mapsto V_y * \mathbf{r} \doteq V_y \\
\{P\}\mathbf{x} = \mathbf{y}\{Q\}
\end{array}$$

Compared to the (Assign Global) and (Assign Local) inference rules, these triples provide an explicit account of the footprint of the assignment, and more clearly describe the destructive effects of assignment. Yet, they are still quite complex and are difficult to compose. It would be useful to ignore some information about the exact structure of the variable store, while retaining the information about the mappings of variable names to values. To do this, we introduce a new store predicate.

6.2 Layer 2: A Simple Abstract Variable Store

The predicates below provide a convenient abstraction for an emulated variable store.

The store Predicate.

$$\begin{array}{l}
\text{store}_L(X_1 \dots X_n | X'_1 : V_1 \dots X'_m : V_m) \triangleq \\
\exists Ls_1 \dots, Ls_n, Ls'_1, \dots, Ls'_m, Ls''_1, \dots, Ls''_m. \text{ thischain}(L) \text{ \# } \\
\text{ \# }_{i \in 1..n} \phi(Ls_i, L, X_i, \text{null}) \text{ \# } (l_g, X_i) \mapsto \emptyset \\
\text{ \# }_{j \in 1..m} \sigma(Ls'_j, L, X'_j, L_j) \text{ \# }_{k \in 1..m} \gamma(Ls''_k, L_k \cdot X'_k, V_k) \\
\text{thischain}([\]) \triangleq \emptyset \\
\text{thischain}(L : Ls') \triangleq (L, @this) \mapsto _ * \text{thischain}(Ls')
\end{array}$$

The assertion $\text{store}_{LS}(\mathbf{a}, \mathbf{b} | \mathbf{x} : 1, \mathbf{y} : 2)$ describes a heap emulating a variable store given by scope chain LS , in which the variables \mathbf{a} and \mathbf{b} are *not* present and the variables \mathbf{x} and \mathbf{y} take the values 1 and 2 respectively. The body of the predicate uses the ϕ predicate to assert the absence of \mathbf{a} and \mathbf{b} , and the σ and γ predicates to assert the presence and values of \mathbf{x} and \mathbf{y} . This information about the exact structure of the emulated store and the locations of \mathbf{x} and \mathbf{y} is hidden from the programmer, since at this level of abstraction it should be of no concern. The *thischain* part of the body of the predicate ensures that the store predicate always encompasses enough resource to call any function that is stored in the emulated variable store. The variables \mathbf{a} and \mathbf{b} can be re-ordered, as can the variables \mathbf{x} and \mathbf{y} . To facilitate program reasoning at this level of abstraction we provide several inference rules, all of which are derived (using previous levels of abstraction) in [11]. We start by giving rules for initialising, overwriting and overriding a variable with a constant and with another variable.

Writing to a store.

$$\begin{array}{l}
\text{Let } Q_1 = \text{store}_e(X_1 \dots X_n | \mathbf{x} : \mathbf{v}, X'_1 : V_1 \dots X'_m : V_m). \\
\text{Let } Q_2 = \text{store}_e(X_1 \dots X_n | \mathbf{x} : \mathbf{v}, \mathbf{y} : V, X'_1 : V_1 \dots X'_m : V_m). \\
\frac{x \neq X_1 \neq \dots \neq X_n \neq X'_1 \neq \dots \neq X'_m}{P = \text{store}_e(\mathbf{x}, X_1 \dots X_n | \mathbf{x} : \mathbf{v}, X'_1 : V_1 \dots X'_m : V_m)} \\
\{P\}\mathbf{x} = \mathbf{v}\{Q_1 * \text{true} * \mathbf{r} \doteq \mathbf{v}\} \\
\frac{x \neq X_1 \neq \dots \neq X_n \neq X'_1 \neq \dots \neq X'_m}{P = \text{store}_e(X_1 \dots X_n | \mathbf{x} : \mathbf{v}, X'_1 : V_1 \dots X'_m : V_m)} \\
\{P\}\mathbf{x} = \mathbf{v}\{Q_1 * \text{true} * \mathbf{r} \doteq \mathbf{v}\} \\
\frac{x \neq \mathbf{y} \neq X_1 \neq \dots \neq X_n \neq X'_1 \neq \dots \neq X'_m}{P = \text{store}_e(\mathbf{x}, X_1 \dots X_n | \mathbf{y} : V, X'_1 : V_1 \dots X'_m : V_m)} \\
\{P\}\mathbf{x} = \mathbf{y}\{Q_2 * \text{true} * \mathbf{r} \doteq \mathbf{v}\}
\end{array}$$

$$\begin{array}{l}
x \neq \mathbf{y} \neq X_1 \neq \dots \neq X_n \neq X'_1 \neq \dots \neq X'_m \\
P = \text{store}_e(X_1 \dots X_n | \mathbf{x} : V', \mathbf{y} : V, X'_1 : V_1 \dots X'_m : V_m) \\
\{P\}\mathbf{x} = \mathbf{y}\{Q_2 * \text{true} * \mathbf{r} \doteq \mathbf{v}\}
\end{array}$$

Notice that each of these rules has a postcondition which includes *true*. This is because initialising or overriding a variable may render a portion of the emulated variable store superfluous. For example, consider the overriding assignment $\mathbf{x} = 1$ in the store given in Figure 1. In this case the cell (l, \mathbf{x}) has become surplus to requirements. In separation logic it is not sound to simply forget about these cells, so we hide them in the general assertion *true*.

One limitation of this level of abstraction is that the abstraction only covers a static (and unknown) scope chain. If we call a function which adds a new cell to the scope chain, then the rules above are insufficient to reason about our program. The following rules allow us to reason at this level of abstraction about a program which alters a global variable from within a new local scope frame.

Writing to a store from a deeper scope.

$$\begin{array}{l}
\text{Let } Q = \text{store}_{LS}(X_1, \dots, X_n | \mathbf{x} : V', X'_1 : V'_1, \dots, X'_m : V'_m) \text{ and } \\
S = (L, @proto) \mapsto \text{null} * (L, \mathbf{x}) \mapsto \emptyset * (L, \mathbf{y}) \mapsto V' * \mathbf{l} \doteq L : LS. \\
\frac{x \neq \mathbf{y} \neq X_1 \neq \dots \neq X_n \neq X'_1 \neq \dots \neq X'_m}{P_1 = \text{store}_{LS}(x, X_1, \dots, X_n | X'_1 : V'_1, \dots, X'_m : V'_m)} \\
\{P_1 * S\}\mathbf{x} = \mathbf{y}\{Q * S * \text{true}\} \\
\frac{x \neq \mathbf{y} \neq X_1 \neq \dots \neq X_n \neq X'_1 \neq \dots \neq X'_m}{P_2 = \text{store}_{LS}(X_1, \dots, X_n | \mathbf{x} : V, X'_1 : V'_1, \dots, X'_m : V'_m)} \\
\{P_2 * S\}\mathbf{x} = \mathbf{y}\{Q * S * \text{true}\}
\end{array}$$

Finally, we provide two rules for a more general case of store interaction. In these cases, the value which is to be written to the variable is the result of computing some arbitrary expression. These specifications are therefore necessarily more complicated, since they must incorporate some features of sequential composition. We insist that whatever the expression does, it must not alter the variable store in a way that changes the visible values of the variables.

Destructive store Initialisation.

$$\begin{array}{l}
\frac{x \neq X_1 \neq \dots \neq X_n \neq X'_1 \neq \dots \neq X'_m}{R = \text{store}_e(\mathbf{x}, X_1 \dots X_n | X'_1 : V_1 \dots X'_m : V_m)} \\
\{R * P\}e\{R \text{ \# } \gamma(LS, V', V) * Q * \mathbf{r} \doteq V'\} \quad \mathbf{r} \notin \text{fv}(Q) \\
S = \left(\frac{\text{store}_e(X_1 \dots X_n | \mathbf{x} : V, X'_1 : V_1 \dots X'_m : V_m)}{\text{ \# } \gamma(LS, V', V) * Q * \text{true} * \mathbf{r} \doteq V} \right) \\
\{R * P\}\mathbf{x} = e\{S\} \\
\frac{x \neq X_1 \neq \dots \neq X_n \neq X'_1 \neq \dots \neq X'_m}{R = \text{store}_{L:SLs}(\mathbf{x}, X_1 \dots X_n | X'_1 : V_1 \dots X'_m : V_m) \text{ \# } \phi(Nsls, \mathbf{l}, \mathbf{x}, L)} \\
\{R\}e\{R \text{ \# } \gamma(LS, V', V) * Q * \mathbf{r} \doteq V'\} \quad \mathbf{r} \notin \text{fv}(Q) \\
S = \left(\frac{\text{store}_{L:SLs}(X_1 \dots X_n | \mathbf{x} : V, X'_1 : V_1 \dots X'_m : V_m)}{\phi(Nsls, \mathbf{l}, \mathbf{x}, L) \text{ \# } \gamma(LS, V', V) * Q * \text{true} * \mathbf{r} \doteq V} \right) \\
\{R * P\}\mathbf{x} = e\{S\}
\end{array}$$

It may seem surprising that we only provide derived rules for destructive variable initialisation, and not for destructive variable *update*. This is because such an update rule would be unsound. The destructive expression might have the side effect of overriding the variable we wish to update. This serves to further demonstrate the need for the low-level reasoning introduced earlier in this paper. We can use high-level abstractions such as the store predicate where they are sound, but if we wish to reason about programs with side-effecting expressions, we will sometimes be forced to reason at a lower level.

The scope of a variable. The store abstraction gives us the tools we need to easily reason about programs with large numbers of variables. For example, consider the program from Section 2:

```

{ storel(x, y, z, f, v) }
x = null; y = null; z = null;
{ storel(f, v | x : null, y : null, z : null) * true }
f = function(w){x=v; v=4; var v; y=v};
{
  { ∃L. storel(v | x : null, y : null, z : null, f : L) *
    (L, @body) → λw. { ... } *
    (L, @scope) → 1 * true }
  v = 5;
  { storel(|x : null, y : null, z : null, f : L, v : 5) *
    (L, @body) → λw. { ... } *
    (L, @scope) → 1 * true }
  f(null);
  { ∃L'. storel(|x : undefined, y : 4, z : null, f : L, v : 5) *
    newobj(L', @proto, @this, w, v) * (L', @proto) → null *
    (L', w) → null * (L', v) → 4 * true }
  [Frame]
  { storel(|x : undefined, y : 4, z : null, f : L, v : 5) }
  z = v;
  { storel(|x : undefined, y : 4, z : 5, f : L, v : 5) * true }
  [Frame]
  { storel(|x : undefined, y : 4, z : 5, f : L, v : 5) *
    newobj(L', @proto, @this, w, v) * (L', @proto) → null *
    (L', w) → null * (L', v) → 4 * true }
  [Cons/Var Elim]
  { ∃L. storel(|x : undefined, y : 4, z : 5, f : L, v : 5) * true }
}

{
  { ∃L', LS. l ≐ L' : LS *
    storeLS(|x : null, y : null, z : null, f : L, v : 5) *
    (L, @body) → λw. { ... } * (L, @scope) → LS *
    newobj(L', @proto, @this, w, v) * (L', @proto) → null *
    (L', w) → null * (L', v) → undefined * (L', @this) → . * true }
  x=v; v=4; var v; y=v;
  { ∃L', LS. l ≐ L' : LS *
    storeLS(|x : undefined, y : 4, z : null, f : L, v : 5) *
    newobj(L', @proto, @this, w, v) * (L', @proto) → null *
    (L', w) → null * (L', v) → 4 * true }
}

```

Figure 2. Reasoning about variable scopes.

```

x = null; y = null; z = null;
f = function(w){x=v;v=4;var v;y=v;};
v = 5; f(null); z = v;

```

With the store predicate and the results given above, reasoning about this program is simple. A proof of the main program is shown in Figure 2. It relies on a simple proof of the function body summarised here and given in full in [11].

Reasoning about with. This level of abstraction also leads itself to reasoning about the notorious with statement. Consider the with example from Section 2 (where f implicitly returns b):

```

a = {b:1}; with (a){f=function(c){b}};
a = {b:2}; f(null)

```

This program demonstrates the importance of modelling with correctly. Recall that when correctly modelled, the closure of the function f will refer to the object $\{b:1\}$, which was pointed to by the variable a at the time that f was defined. However, even though the variable a is changed to point to a different object before $f(\text{null})$ is called, the closure continues to point to the object $\{b:1\}$. Thus the program normally returns the value 1, not 2. Recall also that we must be careful with the f field of the object l_{op} .

We can reason about this program using the store predicate. The proof is in Figure 3. This proof relies on a sub-proof for the invocation of the function $f(\text{null})$, which culminates with the judgement $\{P\}b\{P * r \doteq 1\}$, where P is

```

Let P = (L, b) → 1 * (L, @proto) → lop * true
{ storel(a, f | lop, f) → ∅ * (lop, @proto) → null }
a = {b:1};
{
  { ∃L. storel(f | a : L) → ∅ * (lop, f) → ∅ *
    (lop, @proto) → null * (L, f) → ∅ * P }
  with (a){
    {
      { ∃LS, L. l ≐ L : LS *
        storeLS(f | a : L) → ∅ *
        (lop, @proto) → null * (L, f) → ∅ * P }
      f=function(c){b};
      { ∃LS, L, F. l ≐ L : LS *
        storeLS(|a : L, f : F) → ∅ *
        (lop, @proto) → null * (L, f) → ∅ *
        (F, @body) → λc. {b} * (F, @scope) → L : LS * P }
    }
  };
  {
    { ∃LS, L, F. l ≐ LS *
      storeLS(|a : L, f : F) → ∅ * (lop, f) → ∅ *
      (L, f) → ∅ * (F, @body) → λc. {b} * (F, @scope) → L : LS * P }
    a = {b:2};
    {
      { ∃LS, L, F, L'. l ≐ LS *
        storeLS(|a : L', f : F) → ∅ * (lop, f) → ∅ *
        (L, f) → ∅ * (L', b) → 2 * (L', f) → ∅ *
        (F, @body) → λc. {b} * (F, @scope) → L : LS * P }
      f(null)
      {
        { ∃LS, L, F, L', LOC. l ≐ LS *
          storeLS(|a : L', f : F) → ∅ * (lop, f) → ∅ *
          (L, f) → ∅ * (L', b) → 2 * (L', f) → ∅ *
          (F, @body) → λc. {b} * (F, @scope) → L : LS *
          (LOC, b) → ∅ * (LOC, @proto) → null * P * r ≐ 1 }
        {r ≐ 1 * true}
      }
    }
  }
}

```

Figure 3. Reasoning about with.

$$\left(\begin{array}{l} \exists LS, L, F, L', LOC. l \doteq LOC : L : LS * \\ \text{store}_{LS}(|a : L', f : F) \rightarrow \emptyset * (l_{op}, f) \rightarrow \emptyset * \\ (l_{op}, @proto) \rightarrow \text{null} * \text{true} * \\ (L, b) \rightarrow 1 * (L, f) \rightarrow \emptyset * (L, @proto) \rightarrow l_{op} * \\ (L', b) \rightarrow 2 * (L', f) \rightarrow \emptyset * (L', @proto) \rightarrow l_{op} * \\ (F, @body) \rightarrow \lambda c. \{b\} * (F, @scope) \rightarrow L : LS * \\ (LOC, b) \rightarrow \emptyset * (LOC, @proto) \rightarrow \text{null} \end{array} \right)$$

For space reasons, we only give the reasoning for the case in which neither a nor f are in the variable store. The same techniques in tandem with the disjunction rule can be used to prove the general precondition.

In both the simple and the general case, we must constrain our precondition with the assertion $(l_{op}, f) \rightarrow \emptyset * (l_{op}, @proto) \rightarrow \text{null}$. The requirement for this term may seem surprising. Consider running the above program in a state satisfying $\text{store}_l(a, f) \rightarrow (l_{op}, f) \rightarrow 4$. In this case, when the assignment to f is made, the function pointer will be written to the cell (L, f) , rather than into the global variable store. Since the variable store does not contain a function value for the variable f , the call to $f(\text{null})$ will cause the program to fault. The problem is potentially even worse if (l_{op}, f) contains a function pointer. In this case, the call to $f(\text{null})$ will not fault, but rather will execute whatever code it finds. This kind of unpredictability could lead to very confusing bugs. In the case of a system like Facebook which attempts to isolate user-application code from host-page code, it could even lead to a security flaw.

6.3 Layer 3: a Recursive Abstract Variable Store

While reasoning using the store predicate, it is possible to handle large numbers of assignments and small numbers of function calls. However, for large numbers of function calls, another abstraction is helpful. Consider an *abstract* variable store as a list of lists of variable-value pairs, with the most local scope frame at the head of the outer list. The list $[[x = 4], [y = 5], [x = 6, z = 7]]$ represents a store in which the global scope contains the variables x and z , an intermediate scope adds the variable y , and the

most local scope overrides the variable x . The list elements of variable-value pairs can be represented in our logical expression language as lists containing two elements. For readability, we use the notation $x = v$ above. We define the recursive store predicate $\text{recstore}_L(\text{NoVars}, \text{Store})$ which describes an abstract variable store Store , which does not contain the variables in the list NoVars .

The Recursive recstore predicate.

$$\begin{aligned} & \text{recstore}_L([x_1', \dots, x_m'], [[x_1 = V_1, \dots, x_n = V_n]]) \triangleq \\ & \quad \text{store}_L(x_1', \dots, x_m' | x_1 : V_1, \dots, x_n : V_n) \\ & \text{recstore}_{L:LS}([x_1', \dots, x_m'], ([x_1 = V_1, \dots, x_n = V_n] : Ls')) \triangleq \\ & \quad \text{recstore}_{LS}([x_1', \dots, x_m'], Ls') * (L, @proto) \mapsto \text{null} * L \neq l_g * \\ & \quad *_{i \in 1..m} (L, x_i') \mapsto \emptyset *_{j \in 1..n} (L, x_j) \mapsto V_j * \\ & \quad \text{nones}_L([x_1, \dots, x_n], Ls') \\ & \text{nones}_L(., []) \triangleq \emptyset \\ & \text{nones}_L(Ls, ([x_1 = V_1, \dots, x_n = V_n] : Ls')) \triangleq \\ & \quad *_{i \in 1..n} ((x_i \in Ls) \vee (x_i \notin Ls \wedge (L, x_i) \mapsto \emptyset)) * \\ & \quad \text{nones}_L(x_1 : \dots : x_n : Ls), Ls') \end{aligned}$$

Notice that recstore uses the store predicate to constrain the most global scope frame in the abstract scope list, while being rather more restrictive about more local scope frames. Local scope frames must be emulated by JavaScript objects which contain fields for each variable defined at that scope level. Furthermore, they must assert the absence of fields which are not in the store (the NoVars), or which would otherwise shadow a variable declared in a more global position in the store. This is handled by the nones predicate. Notice that local scope frames have a null prototype, and may not be the l_g object. These criteria are met by the emulated scope frames created by a normal function call, and are not normally met by with calls. This makes this abstraction ideal for reasoning about programs with many function calls and no internal uses of the with statement. Notice, however, that we do not outlaw with calls in the enclosing scope, represented here by a top-level use of the store predicate. This means that this abstraction will facilitate reasoning about libraries which are written in a principled way, and which may be called by unprincipled clients.

We provide several rules for reasoning at this level of abstraction in [11], the most interesting of which are destructive variable initialisation and update.

Destructive recstore update.

$$\begin{aligned} & R = \text{recstore}_1((x : \text{EmpVars}), (\text{Locals} \# [\text{Globals}])) \\ & \{R * P\} \mathbf{e} \{R * Q * \mathbf{r} \doteq \text{Var}\} \\ & \mathbf{r} \notin \text{fv}(Q) \\ & S = \text{recstore}_1((\text{EmpVars}), (\text{Locals} \# [x = \text{Var} : \text{Globals}])) \\ & \{R * P\} \mathbf{x} = \mathbf{e} \{S * Q * \text{true}\} \\ & R = \text{recstore}_1((\text{Emps}), (Ls \# ((x = \text{Var}) : \text{Curr}) \# \text{Globals})) \\ & \{R * \text{Globs} \neq [] * P\} \mathbf{e} \{R * \text{Globs} \neq [] * Q * \mathbf{r} \doteq \text{Var}'\} \\ & \mathbf{r} \notin \text{fv}(Q) \\ & \forall LS \in Ls. (x = _) \notin LS \\ & S = \text{recstore}_1((\text{Emps}), (Ls \# ((x = \text{Var}') : \text{Curr}) \# \text{Globals})) \\ & \{R * \text{Globs} \neq [] * P\} \mathbf{x} = \mathbf{e} \{S * \text{Globs} \neq [] * Q * \mathbf{r} \doteq \text{Var}'\} \end{aligned}$$

Notice that we may not safely update variables in the global portion of the abstract variable store with the results of potentially destructive expressions. This is for the same reason as the corresponding restriction on the store predicate in Section 6.2; there is a corner case which would lead to very unexpected behaviour. At this level of abstraction however, we have an advantage. We can be sure that the local abstract scope frames were constructed in a more principled way, and so we are able to reason about updating them with destructive expressions using the second rule above.

$$\left\{ \begin{array}{l} \text{recstore}_1 \left([], \left[\left[\begin{array}{l} \text{data} = L, \\ \text{checkField} = \&\text{undefined}, \\ i = \&\text{undefined} \end{array} \right], [] \right] \right) * \\ (L, \text{numEntries}) \mapsto N * (L, \text{buttonDisabled}) \mapsto _ * \\ (L, 0) \mapsto \text{TXT}_0 * \dots * (L, N) \mapsto \text{TXT}_N \\ \dots \text{checkForm} \dots \end{array} \right\} \\ \left\{ \begin{array}{l} \exists L'. \text{recstore}_1([], [[\text{data} = L, \text{checkField} = L', i = N], []]) * \\ (L, \text{numEntries}) \mapsto N * \\ (L, 0) \mapsto \text{TXT}_0 * \dots * (L, N) \mapsto \text{TXT}_N * \\ \left(\begin{array}{l} \text{TXT}_0 \neq _ * \dots * \text{TXT}_N \neq _ * \\ (L, \text{buttonDisabled}) \mapsto 0 \end{array} \right) \\ \vee (L, \text{buttonDisabled}) \mapsto 1 \end{array} \right\}$$

Figure 4. The specification of checkForm .

Form validation. Consider a web form with a number of mandatory text fields and a submit button. If the button is “disabled” when the page loads, then an event handler on the form can be used to regularly check if valid data has been entered in all the fields before enabling the button. Let us assume that the programmer has separated the concerns of parsing the web page and of validating the data. The data validation function will be called with a single parameter: an object with one field for each text value to check, a count of those text values, and boolean toggle corresponding to whether the submit button should be disabled. An example function which might perform the validation check is:

```
checkForm = function(data) {
  data.buttonDisabled = 0;
  var checkField = function(text) {
    if(text == "") {data.buttonDisabled = 1;}}
  var i = 0;
  while(i < data.numEntries) {
    checkField(data[i]); i = i+1;}}
```

Notice that this code deals with variables in a principled way. It makes use of no global variables, preferring instead to use function parameters and local variables. The repeated work of the loop body is factored into a function which could be expanded to provide extra functionality or used elsewhere with little cost in readability. Using the recstore abstraction it is straightforward to show that the function body satisfies the specification given in Figure 4.

7. Related Work

We believe this paper is the first to propose a program logic for reasoning about JavaScript. Our program logic adapts ideas from separation logic, and proves soundness with respect to a big-step operational semantics. In this section, we discuss related work on separation logic and the semantics of JavaScript.

We build on the seminal work of O’Hearn, Reynolds and Yang [20], who introduced separation logic for reasoning about C-programs, and on the work of Parkinson and Bierman [4, 23, 24], who adapted separation logic to reason about Java. We made several adaptations to their work in order to reason about JavaScript. As in [4, 23, 24], we use assertions of the form $(1, x) \mapsto 5$ to denote that a field x in object 1 has value 5. We extend these assertions by $(1, x) \mapsto \emptyset$, which denotes that the field is *not* in 1 . This is inspired by Dinsdale-Young *et al.*’s use of the ‘out’ predicate to state that values are not present in a concurrent set [7]. We introduce the sepsch connective \boxtimes to account for partially-shared data structures. We have not seen this connective before, which is surprising since shared data structures are common for example in Linux. There has been much work on various forms of concurrent separation logic with sharing [9, 21, 35], but they all seem to take a different approach to our \boxtimes connective.

We prove the soundness of our frame rule using a new technique developed in Smith’s PhD thesis [30]. There are several approaches to proving soundness of the separation-logic frame rule, and we list the key developments here. The first and most commonly used approach by O’Hearn, Reynolds and Yang involves first proving that the commands of the programming language under consideration are *local*, and then using that property to prove the frame rule [20]. Later, Birkedal and Yang devised a method “baking in” the soundness of the frame rule into their definition of Hoare triples [5]. This made it possible to prove the frame rule without reference to locality, but did not extend well to concurrent programs. Vafeiadis solved the concurrency problem in [34], where he proves soundness of concurrent separation logic without reference to locality. Smith’s thesis was written at the same time as [34], and takes a different approach. Rather than removing all reference to locality from the soundness proof, Smith uses a generalised *weak locality* property, which allows us to handle some non-local behaviour.

We prove our soundness result with respect to a big-step operational semantics of JavaScript, faithfully following the small-step semantics of Maffeis *et al.* [18] except where discussed in Section 3.5. Their semantics captures the complete ECMAScript 3 language, at the same level of abstraction to where a JavaScript programmer reasons. In contrast, Herman and Flanagan [13] provide a definitional interpreter of JavaScript written in ML, which has the advantage of being directly executable, but includes implementation details that obscure the semantic meaning. Elsewhere, Guha *et al.* [12] compile JavaScript to an intermediate Scheme-like language, in order to provide type-based analyses on the object language. They justify their translation in a novel way, by checking that it satisfies an established test suite of JavaScript programs. However, the translation is complicated and it is not immediately apparent that it is sound. In fact, many approaches translate JavaScript programs using troublesome features such as the `with` construct into programs which do not. For example, Park, Lee and Ryu show that this is often feasible [22], but do not handle interactions between `with` and `eval`. There are a number of more abstract models of JavaScript, which have proven useful to study selected language features [1, 32, 37], but which are not sufficiently concrete for our purpose. Overall, we have chosen the semantics in [18] because it appears to be the most faithful to the actual JavaScript semantics. As Richards *et al.* argue in [28], all the unusual features of JavaScript are well-used in the wild, and cannot be easily abstracted away.

8. Conclusions and Future Work

We have defined a program logic for reasoning about JavaScript, based on an operational semantics which is essentially faithful to ECMAScript 3. We have adapted separation logic to reason about a JavaScript subset, modelling many complex features such as prototype inheritance and `with`. We reason about the full dynamic nature of JavaScript’s functions, but do not provide higher-order reasoning. We also provide only conservative reasoning about `eval`. Full reasoning about these features will be technically challenging, although we believe that we can build on the recent work of [6, 10, 29]. We do not currently reason about implicit type coercion. JavaScript allows the programmer to define functions to convert values between different types. Adding this feature to our operational semantics is simple. The complexity in handling it with our program reasoning will be in reasoning about higher-order functions, as discussed above.

Due to our choice of operational semantics, we have been able to prove a strong soundness result. Any library, formally specified using our reasoning, will behave well, even when called by arbitrary JavaScript code. In Section 2, we illustrated this property by demonstrating that our `with` example behaves according to its

specification even when embedded in a program which maliciously uses the non-standard `__proto__` feature. Finally, our soundness result can be extended compositionally to include more sophisticated reasoning about higher-order functions and `eval`.

We have given several examples of our reasoning, demonstrating with small code examples that JavaScript is fiendish to understand and our reasoning can help. In particular, our `with` example shows a potential bug that could easily go unnoticed for some time, and perhaps lead to security holes in sanitised mashup environments. Despite the complexity of the language and the subtlety of the bug, reasoning about this and other examples is made surprisingly simple by our example abstraction layers.

We hope that this work will form the core of a larger body of work on client-side web programming. For example, Thiemann [33] has defined a type-safe DOM API, and Smith [30] has developed context-logic reasoning about DOM Core Level 1. It would be valuable to integrate these approaches to DOM modelling with the JavaScript reasoning presented here.

Maffeis *et al.* [16, 17, 19] have developed techniques for building secure JavaScript mashups out of untrusted code. They prove security properties of their enforcement mechanisms, but do not study their functional correctness. Our reasoning makes such an analysis possible. We also intend to develop reasoning for high-level libraries such as jQuery, Prototype.js and Slidy [15, 25, 26]. We will make each of these libraries the focus of its own layer of abstraction, following the examples in Section 6, to make reasoning about idiomatic uses of those libraries very natural. As with our existing layers of abstraction, it will be possible to safely break these abstractions where necessary in order to code and reason at a lower level.

While this paper deals with ECMAScript 3, the newer ECMAScript 5 and the ongoing ECMAScript Harmony both provide enticing targets. Taly and others have recently given a small-step operational semantics of ECMAScript 5 [31]. It would certainly be interesting to reason about ECMAScript 5 code, and to formalise the connections with our reasoning presented here about ECMAScript 3. We are particularly interested in understanding the behaviour of libraries written in ECMAScript 5 when called by code written in ECMAScript 3.

To make this program reasoning genuinely useful for JavaScript programmers, it is essential that we provide tool support. We intend to produce analysis tools capable of spotting bugs such as the one described in the `with` example in Section 6.2, and integrate our tools with IDEs such as Eclipse or Visual Studio.

Acknowledgements.

Gardner and Smith are supported by EPSRC grant COTFM.P21654. Maffeis is supported by EPSRC grant EP/I004246/1. We thank Charlton, Dinsdale-Young, Ntzik, O’Hearn, da Rocha Pinto, Wheelhouse, and Wright for insightful discussions. We especially thank Daiva Naudziuniene for her impressive eye for detail.

References

- [1] C. Anderson, P. Giannini, and S. Drossopoulou. Towards type inference for JavaScript. In *Proc. of ECOOP’05*, 2005. 7
- [2] J. Berdine, C. Calcagno, and P. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *FMCO*, 2005. 1
- [3] J. Berdine, B. Cook, and S. Ishtiaq. Slayer: Memory safety for systems-level code. In *CAV*, 2011. 1
- [4] G.M. Bierman, M.J. Parkinson, and A. M. Pitts. MJ: An imperative core calculus for java and java with effects. Technical report, Cambridge, 2003. 4, 7
- [5] L. Birkedal and H. Yang. Relational parametricity and separation logic. In *FoSSaCS*, pages 93–107, 2007. 7

- [6] N. Charlton. Hoare logic for higher order store using simple semantics. In *Proc. of WOLLIC 2011*, 2011. 1, 8
- [7] T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. *ECOOP*, 2010. 2, 7
- [8] D. Distefano and M. Parkinson. jStar: towards practical verification for Java. In *OOPSLA '08*, pages 213–226. ACM, 2008. 1
- [9] M. Dodds, X. Feng, M.J. Parkinson, and V. Vafeiadis. Deny-guarantee reasoning, 2009. 7
- [10] D. Dreyer, G. Neis, and L. Birkedal. The impact of higher-order state and control effects on local relational reasoning. In *ICFP*, pages 143–156, 2010. 8
- [11] P. Gardner, S. Maffeis, and G. Smith. Towards a program logic for JavaScript. Imperial College London Technical Report number DTR11-11, November 2011. 3.1, 3.3, 3.3, 3.4, 4, 5.1, 5.2, 5.3, 6.1, 6.2, 6.3
- [12] A. Guha, C. Saftoiu, and S. Krishnamurthi. The Essence of JavaScript. *ECOOP 2010*, pages 126–150, 2010. 1, 2, 7
- [13] D. Herman and C. Flanagan. Status report: specifying JavaScript with ML. In *Proc. of ML'07*, pages 47–52, 2007. 7
- [14] ECMA International. ECMAScript language specification. standard ECMA-262, 3rd Edition, 1999. 1, 2
- [15] jQuery: The write less, do more, JavaScript library. <http://jquery.com>. 8
- [16] S. Maffeis, J. C. Mitchell, and A. Taly. Isolating javascript with filters, rewriting, and wrappers. In *ESORICS*, pages 505–522, 2009. 8
- [17] S. Maffeis, J. C. Mitchell, and A. Taly. Object capabilities and isolation of untrusted web applications. In *IEEE Symposium on Security and Privacy*, pages 125–140, 2010. 8
- [18] S. Maffeis, J.C. Mitchell, and A. Taly. An operational semantics for JavaScript. In *Proc. of APLAS'08*, LNCS, 2008. 1, 2, 3, 3.2, 7
- [19] S. Maffeis and A. Taly. Language-based isolation of untrusted javascript. In *CSF*, pages 77–91, 2009. 8
- [20] P. O'Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *CSL*, 2001. 1, 4, 5.3, 7
- [21] P. W. O'Hearn. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.*, 375(1-3):271–307, 2007. 7
- [22] Changhee Park, Hongki Lee, and Sukyoung Ryu. An empirical study on the rewritability of the with statement in javascript. In *FOOL*, 2011.
- [23] M. J. Parkinson and G. M. Bierman. Separation logic, abstraction and inheritance. In *POPL*, 2008. 2, 7
- [24] M. J. Parkinson. Local reasoning for Java. Technical Report 654, Univ. of Cambridge Computer Laboratory, 2005. Ph.D. dissertation. 4, 7
- [25] Prototype Core Team. Prototype JavaScript framework: Easy Ajax and DOM manipulation for dynamic web applications. <http://www.prototypejs.org>. 4, 7
- [26] Dave Raggett. W3C Slidy. <http://www.w3.org/Talks/Tools/Slidy2/>, 2005. 8
- [27] G. Richards, C. Hammer, B. Burg, and J. Vitek. The Eval that men do A large-scale study of the use of Eval in JavaScript applications. Accepted for publication at ECOOP 2011. 8
- [28] G. Richards, S. Lebesne, B. Burg, and J. Vitek. An analysis of the dynamic behavior of JavaScript programs. In *PLDI*, 2010. 2
- [29] J. Schwinghammer, L. Birkedal, B. Reus, and H. Yang. Nested hoare triples and frame rules for higher-order store. In *In Proc. of CSL'09*, 2009. 1, 2, 2, 7
- [30] G. D. Smith. Local reasoning about web programs. PhD Thesis, Dep. of Computing, Imperial College London, 2011. 1, 8
- [31] A. Taly, U. Erlingsson, M. S. Miller, J. C. Mitchell, and J. Nagra. Automated analysis of security-critical javascript apis. In *Proc. of IEEE Security and Privacy'11*. IEEE, 2011. 1, 5.3, 7, 8
- [32] P. Thiemann. Towards a type system for analyzing javascript programs. In *Proc. of ESOP'05*, volume 3444 of LNCS, 2005. 8
- [33] P. Thiemann. A type safe DOM API. In *Proc. of DBPL*, pages 169–183, 2005. 7
- [34] V. Vafeiadis. Concurrent separation logic and operational semantics. In *MFPS11*, 2011. 8
- [35] Viktor Vafeiadis and M. Parkinson. A marriage of rely/guarantee and separation logic. In *IN 18TH CONCUR*. Springer, 2007. 7
- [36] H. Yang, O. Lee, J. Berdine, C. Calcagno, B. Cook, D. Distefano, and P. O'Hearn. Scalable shape analysis for systems code. In *CAV*, 2008. 7
- [37] D. Yu, A. Chander, N. Islam, and I. Serikov. JavaScript instrumentation for browser security. In *Proc. of POPL'07*, 2007. 1