An Operational Semantics for JavaScript[*]

Sergio Maffeis[1], John C. Mitchell[2], Ankur Taly[2],

[1] Department of Computing, Imperial College London
[2] Department of Computer Science, Stanford University

Abstract. We define a small-step operational semantics for the ECMAScript standard language corresponding to JavaScript, as a basis for analyzing security properties of web applications and mashups. The semantics is based on the language standard and a number of experiments with different implementations and browsers. Some basic properties of the semantics are proved, including a soundness theorem and a characterization of the reachable portion of the heap.

1 Introduction

JavaScript [8,15,10] is widely used in Web programming and it is implemented in every major browser. As a programming language, JavaScript supports functional programming with anonymous functions, which are widely used to handle browser events such as mouse clicks. JavaScript also has objects that may be constructed as the result of function calls, without classes. The *properties* of an object, which may represent methods or fields, can be inherited from a prototype, or redefined or even removed after the object has been created. For these and other reasons, formalizing JavaScript and proving correctness or security properties of JavaScript mechanisms poses substantial challenges.

Although there have been scientific studies of limited subsets of the language [7,24,27], there appears to be no previous formal investigation of the full core language, on the scale defined by the informal ECMA specifications [15]. In order to later analyze the correctness of language-based isolation mechanisms for JavaScript, such as those that have arisen recently in connection with online advertising and social networking [1,2,6,23], we develop a small-step operational semantics for JavaScript that covers the language addressed in the ECMA-262 Standard, 3rd Edition [15]. This standard is intended to define the common core language implemented in all browsers and is roughly a subset of JavaScript 1.5. We provide a basis for further analysis by proving some properties of the semantics, such as a progress theorem and properties of heap reachability.

As part of our effort to make conformance to the informal standard evident, we define our semantics in a way that is faithful to the common explanations of JavaScript and the intuitions of JavaScript programmers. For example, JavaScript scope is normally discussed in relation to an object-based representation. We therefore define execution of a program with respect to a heap that contains a linked structure of objects instead of a separate stack. Thus entering a JavaScript scope creates an object on the heap, serving as an activation record for that scope but also subject to additional operations on JavaScript objects. Another unusual aspect of our semantics, reflecting the

---

[*] This is a revised and extended version of the conference paper [17] that appeared in the Proceedings of APLAS 2008.

unusual nature of JavaScript, is that declarations within the body of a function are handled by a two-pass method. The body of a function is analyzed for declarations, which are then added to the scope before the function body is executed. This allows a declaration that appears after the first expression in the function body to be referenced in that expression.

While the ECMAScript language specification guided the development of our operational semantics, we performed many experiments to check our understanding of the specification and to determine differences between various implementations of JavaScript. The implementations that we considered include SpiderMonkey [19] (used by Firefox), the Rhino [4] implementation for Java, JScript [3] (used by Internet Explorer), and the implementations provided in Safari and Opera. In the process, we developed a set of programs that test implementations against the standard and reveal details of these implementations. Many of these program examples, a few of which appear further below, may be surprising to those familiar with more traditional programming languages. Because of the complexity of JavaScript and the number of language variations, our operational semantics (reported in full in [16]) is approximately 70 pages of rules and definitions, in ascii format. We therefore describe only a few of the features and implications of the semantics here. By design, our operational semantics is modular in a way that allows individual clauses to be varied to capture differences between implementations.

Since JavaScript is an unusual language, there is value and challenge in proving properties that might be more straightforward to verify for some other languages (or for simpler idealized subsets of JavaScript). We start by proving a form of soundness theorem, stating that evaluation progresses to an exception or a value of an expected form. Our second main theorem shows, in effect, that the behavior of a program depends only on a portion of the heap. A corollary is that certain forms of garbage collection, respecting the precise characterization of heap reachability used in the theorem, are sound for JavaScript. This is non-trivial because JavaScript provides a number of ways for an expression to access, for example, the parent of a parent of an object, or even its own scope object, increasing the set of potentially reachable objects. The precise statement of the theorem is that the operational semantics preserve a similarity relation on states (which include the heap).

There are several reasons why the reachability theorem is important for various forms of JavaScript analysis. For example, a web server may send untrusted code (such as an advertisement) as part of a trusted page (the page that contains third-party advertisement). We would therefore like to prove that untrusted code cannot access certain browser data structures associated with the trusted enclosing page, under specific conditions that could be enforced by web security mechanisms. This problem can be reduced to proving that a given well-formed JavaScript program cannot access certain portions of the heap, according to the operational semantics of the language. Another future application of heap bisimilarity (as shown in this paper) to security properties of JavaScript applications is that in the analysis of automated phishing defenses, we can reduce the question of whether JavaScript can distinguish between the original page and a phishing page to whether there exists a bisimulation between a certain good heap (corresponding to the original page) and a certain bad heap (corresponding to the phishing page). Thus the framework that we develop in this paper for proving basic progress and heap reachability theorems provides a useful starting point for JavaScript security mechanisms and their correctness proofs.

1.1 JavaScript Overview and Challenges

JavaScript was originally designed to be a simple HTML scripting language [8]. The main primitives are first-class and potentially higher-order functions, and a form of object that can be defined by an object expression, without the need for class declarations. Commonly, related objects are constructed by calling a function that creates objects and returns them as a result of the function call. Functions and objects have *properties,* which are accessed via the "dot" notation, as in x.p for property p of object x. Properties can be added to an object or reset by assignment. This makes it conceptually possible to represent activation records by objects, with assignable variables considered properties of the object corresponding to the current scope. Because it is possible to change the value of a property arbitrarily, or remove it from the object, static typing for full JavaScript is difficult. JavaScript also has eval, which can be used to parse and evaluate a string as an expression, and the ability to iterate over properties of an object or access them using string expressions instead of literals (as in x["p"]). Many online tutorials and books [10] are available.

One example feature of JavaScript that is different from other languages that have been formally analyzed is the way that declarations are processed in an initial pass before bytecode for a function or other construct is executed. Some details of this phenomenon are illustrated by the following code:

```
var f = function(){if (true) {function g() {return 1}}
                   else {function g() {return 2}};
                function g() {return 3};
                return g();
                function g() {return 4}}
```

This code defines a function f whose behavior is given by one of the declarations of g inside the body of the anonymous function that returns g. However, different implementations disagree on which declaration determines the behavior of f. Specifically, a call f() should return 4 according to the ECMA specification. Spidermonkey (hence Firefox) returns 4, while Rhino and Safari return 1, and JScript and the ECMA4 reference implementation return 2. Intuitively, the function body is parsed to find and process all declarations before it is executed, so that reachability of second declarations is ignored. Given that, it is plausible that most implementations would pick either the first declaration or the last. However, this code is likely to be unintuitive to most programmers.

Those with some curiosity may enjoy the following example on the difference between a declaration function f(...){ ... } and the apparently semantically similar var f = function (...){ ... } which uses another form of binding to associate the same name with an apparently equivalent function.

```
function f(x){ if ( x == 0){return 1;}else{return f(x−1);}};
var h = f;
h(3) % res: 1
function f(x){ if ( x == 0){return 3;}else{return x∗f(x−1);}};
h(3) % res: 6
```

Unsurprisingly, the call to h(3) after the second line evaluates to 1. However, the call to h(3) after the third line produces 6. In effect, the call to h(3) first executes the first body of f, apparently because that's the declaration of f that was current at the place

where h was declared. However, the recursive call to f in the body of line one invokes the declaration on the third line! This and other examples lead us to suggest the var f = function (...){ ...} form for programming, since this form avoids various anomalies. However, the semantics accurately treats both forms.

Additional Challenges. A number of features of JavaScript provide additional challenges for development of a formal semantics and proving properties of the language. We list some of them:

- *Redefinition.* Values undefined, NaN and Infinity, but especially Object, Function and so on can be redefined. Therefore the semantics cannot depend on fixed meanings for these predefined parts of the language.
- *Implicit mutable state.* Some JavaScript objects, such as Array.prototype, are implicitly reachable even without naming any variables in the global scope. The mutability of these objects allows apparently unrelated code to interact.
- *Constrained Properties.* Some properties of native JavaScript objects are constrained to be Internal, ReadOnly, DontEnum, or DontDelete, but there is no mechanism to express these constraints in the (client) language.
- *Property Enumeration.* JavaScript's for/in loop enumerates the properties of an object, whether inherited or not, unless the property is DontEnum.
- *Enumeration order.* The ECMA standard [15] does not define the order of enumeration of properties in a for/in loop, leading to divergent implementations.
- *"this" confusion.* JavaScript's rules for binding this depend on whether a function is invoked as a constructor, as a method, as a normal function, etc.. If a function written to be called in one way is instead called in another way, its this property might be bound to an unexpected object or even to the global environment.

1.2 Beyond this Paper

Our framework for studying the formal properties of JavaScript closely follows the specification document and models all the features of the language that we have considered necessary to represent faithfully its semantics. The semantics can be modularly extended to *user-defined getters and setters*, which are part of JavaScript 1.5 but not of the ECMA-262 standard. We believe it is similarly possible to extend the semantics to interface with *DOM objects*, which are part of an independent specification (a formal subset is presented in [11]), and are available only when JavaScript runs in a Web-browser. However, we leave development of these extensions to future work.

For simplicity, we do not model some features which are laborious but do not add new insight to the semantics, such as the switch and for construct (we do model the for−in), parsing (which is used at run time for example by the eval command), the native Date and Math objects, minor type conversions like ToUInt32, etc. and the details of standard procedures such as converting a string into the numerical value that it actually represents. For the same reason, we also do not model *regular expression matching*, which is used in string operations.

In Section 6 we summarize some directions for future work.

2 Operational Semantics

Our operational semantics consists of a set of rules written in a conventional meta-notation. The notation is not directly executable in any specific automated framework,

but is designed to be humanly readable, insofar as is possible for a programming language whose syntax requires 16 pages of specification, and a suitable basis for rigorous but un-automated proofs. Given the space constraints of a conference paper, we describe only the main semantic functions and some representative axioms and rules; the full semantics is currently available online [16].

In order to keep the semantic rules concise, we assume that source programs are legal JavaScript programs, and that each expression is disambiguated (e.g. 5+(3∗4)). We also follow systematic conventions about the syntactic categories of metavariables, to give as much information as possible about the intended type of each operation. In addition to the source expressions and commands used by JavaScript programmers, our semantics uses auxiliary syntactic forms that conveniently represent intermediate steps in our small-step semantics.

In principle, for languages whose semantics are well understood, it may be possible to give a direct operational semantics for a core language subset, and then define the semantics of additional language constructs by showing how these additional constructs are expressible in the core language. Instead of assuming that we know how to correctly define some parts of JavaScript from others, we decided to follow the ECMA specification as closely as possible, defining the semantics of each construct directly as given in the ECMA specification. While giving us the greatest likelihood that the semantics is correct, this approach also did not allow us to factor the language into independent sublanguages. While our presentation is divided into sections for *Types*, *Expressions*, *Objects*, and so on, the execution of a program containing one kind of construct may rely on the semantics of other constructs. We consider it an important future task to streamline the operational semantics and prove that the result is equivalent to the form derived from the standard.

Syntactic Conventions. In the rest of the paper we abbreviate t1,..., tn with t˜ and t1 ... tn with t∗ (t+ in the nonempty case). In a grammar, [t] means that t is optional, t|s means either t or s, and in case of ambiguity we escape with apices, as in escaping [by "[". Internal values are prefixed with &, as in &NaN. For conciseness, we use short sequences of letters to denote metavariables of a specific type. For example, m ranges over strings, pv over primitive values, etc.. These conventions are summarized in Figure 1. In the examples, unless specified otherwise, JavaScript code prefixed by *js>* is verbatim code from the SpiderMonkey shell (release 1.7.0 2007-10-03).

2.1 Heap

Heaps and Values. Heaps map locations to objects, which are records of pure values va or functions fun(x,...){P}, indexed by strings m or internal identifiers @x (the symbol @ distinguishes internal from user identifiers). Values are standard. As a convention, we append w to a syntactic category to denote that the corresponding term may belong to that category or be an exception. For example, lw denotes an address or an exception.

Heap Functions. We assume a standard set of functions to manipulate heaps. alloc(H,o) = H1,l allocates o in H returning a fresh address l for o in H1. H(l) = o retrieves o from l in H. o.i = va gets the value of property i of o. o−i = fun([x˜]){P} gets the function stored in property i of o. o:i = {[a˜]} gets the possibly empty set of attributes of property i of o. H(l.i=ov)=H1 sets the property i of l in H to the object value ov. del(H,l,i) = H1 deletes i from l in H. i !< o holds if o does not have property i. i < o holds if o has property i.

```
H  ::= (l:o)˜ % heap
l  ::= #x % object addresses
x  ::= foo | bar | ... % identifiers (do not include reserved words)
o  ::= ″{″[(i:ov)˜]″}″ % objects
i  ::= m | @x % indexes
ov ::= va[″{″a˜″}″] % object values
       | fun″(″[x˜]″){″P″}″ % function
a  ::= ReadOnly| DontEnum | DontDelete % attributes

pv ::= m | n | b | null | &undefined % primitive values
m  ::= ″foo″ | ″bar″ | ... % strings
n  ::= −n | &NaN | &Infinity | 0 | 1 | ... % numbers
b  ::= true | false % booleans
va ::= pv | l % pure values
r  ::= ln″∗″m % references
ln ::= l | null % nullable addresses
v  ::= va | r % values
w  ::= ″<″va″>″ % exception
```

Fig. 1. Metavariables and Syntax for Values

2.2 Semantics Functions

We have three small-step semantic relations for expressions, statements and programs denoted respectively by $\xrightarrow{e}, \xrightarrow{s}, \xrightarrow{P}$. Each semantic function transforms a heap, a pointer in the heap to the current scope, and the current term being evaluated into a new heap-scope-term triple. The evaluation of expressions returns either a value or an exception, the evaluation of statements and programs terminates with a completion (explained below).

The semantic functions are recursive, and mutually dependent. The semantics of programs depends on the semantics of statements which in turn depends on the semantics of expressions which in turn, for example by evaluating a function, depends circularly on the semantics of programs. These dependencies are made explicit by contextual rules, that specify how a transition derived for a term can be used to derive a transition for a bigger term including the former as a sub-term. In general, the premises of each semantic rule are predicates that must hold in order for the rule to be applied, usually built of very simple mathematical conditions such as $t < S$ or $t \mathrel{!=} t'$ or $f(a) = b$ for set membership, inequality and function application.

Transitions axioms (rules that do not have transitions in the premises) specify the individual transitions for basic terms (the redexes). The axiom $H,l,(v) \longrightarrow H,l,v$, for example, describes that brackets can be removed when they surround a value (as opposed to an expression, where brackets are still meaningful).

Contextual rules propagate such atomic transitions. For example, if program H,l,P evaluates to $H1,l1,P1$ then $H,l,@FunExe(l',P)$ (an internal expression used to evaluate the body of a function) reduces in one step to $H1,l1,@FunExe(l',P1)$. The rule below show

exactly that: @FunExe(l,−) is one of the contexts eCp for evaluating programs.

$$\frac{\text{H,l,P} \xrightarrow{P} \text{H1,l1,P1}}{\text{H,l,eCp[P]} \xrightarrow{e} \text{H1,l1,eCp[P1]}}$$

As another example, sub-expressions are evaluated inside outer expressions (rule on the left) using contexts eC ::= typeof eC | eCgv | ..., and exceptions propagated to the top level (axiom on the right).

$$\frac{\text{H,l,e} \xrightarrow{e} \text{H1,l1,e1}}{\text{H,l,eC[e]} \xrightarrow{e} \text{H1,l1,eC[e1]}} \qquad \text{H,l,eC[w]} \xrightarrow{e} \text{H,l,w}$$

Hence, if an expression throws an exception (H,l,e $\xrightarrow{e}$ H1,l,w) then so does say the typeof operator: H,l,typeof e $\xrightarrow{e}$ H1,l,typeof w $\xrightarrow{e}$ H1,l,w.

It is very convenient to nest contexts inside each other. For example, contexts for GetValue (the internal expression that returns the value of a reference), generated by the grammar for eCgv ::= −[e] | va[−] | eCto | eCts | ..., are expression contexts. Similarly, contexts for converting values to objects eCto ::= −[va] | ... or to strings eCts ::= l[−] | ... are get-value contexts.

$$\text{H,l,eCgv[ln∗m]} \xrightarrow{e} \text{H1,l,eCgv[@GetValue(ln∗m)]}$$

$$\frac{\text{Type(va)} \mathrel{!=} \text{Object} \quad \text{ToObject(H,va)} = \text{H1,lw}}{\text{H,l,eCto[va]} \xrightarrow{e} \text{H1,l,eCto[lw]}}$$

$$\frac{\text{Type(v)} \mathrel{!=} \text{String} \quad \text{ToString(v)} = \text{e}}{\text{H,l,eCts[v]} \xrightarrow{e} \text{H,l,eCts[e]}}$$

As a way to familiarize with these structures of nested contexts, we look in detail at the concrete example of member selection. The ECMA-262 specification states that in order to evaluate `MemberExpression[Expression]` one needs to:

```
MemberExpression : MemberExpression [ Expression ]
1. Evaluate MemberExpression.
2. Call GetValue(Result(1)).
3. Evaluate Expression.
4. Call GetValue(Result(3)).
5. Call ToObject(Result(2)).
6. Call ToString(Result(4)).
7. Return a value of type Reference whose base object is Result(5) and
whose property name is Result(6).
```

In our formalization, the rule for member selection is just H,l,l1[m] $\xrightarrow{e}$ H,l,l1∗m. As opposed to the textual specification, the formal rule is trivial, and makes it obvious that the operator takes an object and a string and returns the corresponding reference. All we had to do in order to model the intermediate steps was to insert the appropriate contexts in the evaluation contexts for expressions, values, objects and strings. In particular, −[e] is value context, and value contexts are expression contexts, so we get for free steps 1 and 2, obtaining va[e]. Since va[−] is also a value context, steps 3 and 4 also come for free, obtaining va[va]. Since −[va] is an object context, and l[−] a string context, steps 6 and 7 are also executed transparently. The return type of each of those

contexts guarantees that the operations are executed in the correct order. If that was not the case, then the original specification would have been ambiguous. The rule for propagating exceptions takes care of any exception raised during these steps.

The full formal semantics [16] contains several other contextual rules to account for other mutual dependencies and for all the implicit type conversions. This substantial use of contextual rules greatly simplifies the semantics and will be very useful in Section 4 to prove its formal properties.

Scope and Prototype Lookup. The scope and prototype chains are two distinctive features of JavaScript. The stack is represented implicitly, by maintaining a chain of objects whose properties represent the binding of local variables in the scope. Since we are not concerned with performance, our semantics needs to know only a pointer to the head of the chain (the current scope object). Each scope object stores a pointer to its enclosing scope object in an internal @Scope property. This helps in dealing with constructs that modify the scope chain, such as function calls and the with statement.

JavaScript follows a prototype-based approach to inheritance. Each object stores in an internal property @Prototype a pointer to its prototype object, and inherits its properties. At the root of the prototype tree there is @Object.prototype, that has a null prototype. The rules below illustrate prototype chain lookup.

$$\text{Prototype}(H,\text{null},m)=\text{null}$$

$$\frac{m < H(l)}{\text{Prototype}(H,l,m)=l} \qquad \frac{m! < H(l) \quad H(l).@\text{Prototype}=ln}{\text{Prototype}(H,l,m)=\text{Prototype}(H,ln,m)}$$

Function Scope(H,l,m) returns the address of the scope object in H that first defines property m, starting from the current scope l. It is used to look up identifiers in the semantics of expressions. Its definition is similar to the one for prototype, except that the condition (H,l.@HasProperty(m)) (which navigates the prototype chain to check if l has property m) is used instead of the direct check $m < H(l)$.

Types. JavaScript values are dynamically typed. The internal types are:

T ::= Undefined | Null | Boolean | String | Number *% primitive types*
    | Object | Reference *% other types*

Types are used to determine conditions under which certain semantic rules can be evaluated. The semantics defines straightforward predicates and functions which perform useful checks on the type of values. For example, IsPrim(v) holds when v is a value of a primitive type, and GetType(H,v) returns a string corresponding to a more intuitive type for v in H. The user expression typeof e, which returns the type of its operand, uses internally GetType. Below, we show the case for function objects (i.e. objects which implement the internal @Call property).

$$\frac{\text{Type}(v)=\text{Object} \quad @\text{Call} < H(v)}{\text{GetType}(H,v) = \text{"function"}}$$

An important use of types is to convert the operands of typed operations and throw exceptions when the conversion fails. There are implicit conversions into strings, booleans, number, objects and primitive types, and some of them can lead to the execution of arbitrary code. For example, the member selection expression e1[e2] implicitly converts e2 to a string. If e2 denotes an object, its re-definable toString propery is invoked as a function.

js> var o={a:0}; o[{toString:function(){o.a++; return *"a"*}}] *% res: 1*

The case for ToPrimitive (invoked by ToNumber and ToString) is responsible for the side effects: the result of converting an object value into a primitive value is an expression l.@DefaultValue([T]) which may involve executing arbitrary code that a programmer can store in the valueOf or toString methods of said object.

$$\frac{\mathsf{Type(l)=Object}}{\mathsf{ToPrimitive(l[,T]) = l.@DefaultValue([T])}}$$

2.3 Expressions

We distinguish two classes of expressions: internal expressions, which correspond to specification artifacts needed to model the intended behavior of user expressions, and user expressions, which are part of the user syntax of JavaScript. The syntax for user expressions is reported in Figure 2, where we use &PO,&UN,&BIN to range respectively over primitive, unary and binary operators.

---

e ::=
    this *% the "this" object*
    x *% identifier*
    pv *% primitive value*
    "[" [e˜] *"]"* *% array literal*
    "{"[(pn:e)˜]*"}"* *% object literal*
    "("e")" *% parenthesis expression*
    e.x *% property accessor*
    e*"["e"]"* *% member selector*
    new e[*"("("e˜")")"*] *% constructor invocation*
    e*"("("e˜")")"* *% function invocation*
    function [x] *"("*[x˜]*")"*{*"["P"]"*}*"* *% [named] function expression*
    e &PO *% postfix operator*
    &UN e *% unary operators*
    e &BIN e *% binary operators*
    "("e"?"e":"e")" *% conditional expression*
    (e,e) *% sequential expression*

pn ::= n | m | x *% property name*

Fig. 2. Syntax for Expressions

Internal expressions include addresses, references, exceptions and functions such as @GetValue,@PutValue used to get or set object properties, and @Call,@Construct used to call functions or to construct new objects using constructor functions. For example, we give two rules of the specification of @Put, which is the internal interface (used also by @PutValue) to set properties of objects. The predicate H,l1.@CanPut(m) holds if m does not have a ReadOnly attribute.

$$\frac{\mathsf{H,l1.@CanPut(m)} \quad\quad \mathsf{m\ !<\ H(l1)} \quad \mathsf{H(l1.m=va\{\})=H1}}{\mathsf{H,l,l1.@Put(m,va)} \xrightarrow{e} \mathsf{H1,l,va}}$$

$$\frac{H,l1.@CanPut(m) \quad H(l1):m=\{[a\tilde{\ }]\} \quad H(l1.m=va\{[a\tilde{\ }]\}) = H1}{H,l,l1.@Put(m,va) \xrightarrow{e} H1,l,va}$$

These rules show that fresh properties are added with an empty set of attributes, whereas existing properties are replaced maintaining the same set of attributes.

Object Literal. As an example of expressions semantics we present in detail the case of object literals. The semantics of the object literal expression {pn:e,...,pn':e'} uses an auxiliary internal construct AddProps to add the result of evaluating each e as a property with name pn to a newly created empty object. Rule (1) (with help from the contextual rules) creates a new empty object, and passes control to AddProps. Rule (2) converts identifiers to strings, and rule (3) adds a property to the object being initialized. It uses a sequential expression to perform the update and then return the pointer to the updated object l1, which rule (4) releases at the top level.

$$H,l,\{[(pn:e)\tilde{\ }]\} \xrightarrow{e} H,l,@AddProps(new\ Object()[,(pn:e)\tilde{\ }]) \quad (1)$$

$$H,l,@AddProps(l1,x:e[,(pn:e)\tilde{\ }]) \xrightarrow{e} H,l,@AddProps(l1,"x":e[,\ (pn:e)\tilde{\ }]) \quad (2)$$

$$H,l,@AddProps(l1,m:va[,(pn:e)\tilde{\ }]) \xrightarrow{e} H,l,@AddProps((l1.@Put(m,va),l1)[,(pn:e)\tilde{\ }]) \quad (3)$$

$$H,l,@AddProps(l1) \xrightarrow{e} H,l,l1 \quad (4)$$

Rule (1) is emblematic of a few other cases in which the specification requires to create a new object by evaluating a specific constructor expression whose definition can be changed during execution. For example,

```
js> var a = {}; a % res: [object Object]
ljs> Object = function(){return new Number()}; var b = {}; b % res: 0
```

where the second object literal returns a number object. That feature can be useful, but can also lead to undesired confusion.

The in operator. By virtue of the contextual rules, the construct e1 in e2 evaluates e2 to a pure value va. If va is an object l1, it evaluates e1 to a string m and returns a boolean b representing whether or not m denotes a property in the prototype chain of l1 (H,l1.@HasProperty(m)=b). If va denotes a primitive value pv, a TypeError exception is thrown.

$$\frac{Type(l1) = Objec\ t\ H,l1.@HasProperty(m)=b}{H,l,m\ in\ l1 \xrightarrow{e} H,l,b}$$

$$\frac{Type(pv) \mathrel{!=} Object \quad o = new\_TypeError \quad H1,l1=alloc(H,o)}{H,l,va\ in\ pv \xrightarrow{e} H1,l,<l1>}$$

Binary logical operators. One may expect binary logical operators to have a dull semantics. That is not the case in JavaScript, as show in the example below.

```
js> var a = true; if (a*1) 1; else 0 % res: 1
js> var b = {valueOf:function(){return 0}}
js> var a = b&&b; if(a*1) 1; else 0 % res: 0
js> if (b) 1; else 0 % res: 1
```

From the semantics it is immediately apparent that the && operator does not return a boolean value, but instead its second argument if the first one evaluates to true, or else its first argument. Combining this with implicit type conversions can be very confusing for a programmer. Below, (^) denotes xor:

$$H,l,va \ \&\& \ e \ \xrightarrow{\ e\ } H,l,@L(true,va,va,e)$$

$$H,l,va \ || \ e \ \xrightarrow{\ e\ } H,l,@L(false,va,va,e)$$

$$\frac{b1(\hat{\ })b2}{H,l,@L(b1,b2,va,e) \ \xrightarrow{\ e\ } H,l,va} \qquad \frac{!(b1(\hat{\ })b2)}{H,l,@L(b1,b2,va,e) \ \xrightarrow{\ e\ } H,l,@GV(e)}$$

Both && and || are translated to the same internal @L expression modulo a flag that is used (in xor with the result of converting the first parameter to a boolean) to determine which parameter to return.

Relational operators. Also relational operators do not fail to surprise. Due to implicit type conversions, the code below has the side effect of setting x to 2. Implementations diverge on this point, so we show a snapshot from the Rhino shell (see Section 3) which respects the ECMA specification.

js> x = 0; var a = {valueOf:function (){x=1}};
    var b = {valueOf:function (){x=2}}; a<b; x *% res: 2*

Yet, the evaluation order is not always left to right!

js> x = 0; var a = {valueOf:function (){x=1}};
    var b = {valueOf:function (){x=2}}; a>b; x *% res: 1*

The reason is that both >,< are translated in terms of the same internal expression (which is a context for converting objects to primitive values), modulo the order of their parameters.

$$H,l,va1<va2 \ \xrightarrow{\ e\ } H,l,@<S(false,va1::Number,va2::Number)$$

$$H,l,va1>va2 \ \xrightarrow{\ e\ } H,l,@<S(false,va2::Number,va1::Number)$$

Equality checks. There are two kinds of equality testing, with or without implicit type conversions. Inequalities are defined as evaluating the logical negation of the corresponding equality.

$$H,l,e1 \ !== \ e2 \ \xrightarrow{\ e\ } H,l,!(e1 === e2)$$

$$H,l,e1 \ != \ e2 \ \xrightarrow{\ e\ } !(e1 == e2)$$

Strict equality can be implemented as loose equality modulo type equality:

$$\frac{Type(va1) = Type(va2)}{H,l,va1===va2 \ \xrightarrow{\ e\ } H,l,va1==va2} \qquad \frac{Type(va1) \ != \ Type(va2)}{H,l,va1===va2 \ \xrightarrow{\ e\ } H,l,false}$$

Loose equality requires lots of different rules to account for the different conversions required by the operands. We give only a few assorted cases, as an example.

$$H,l,null==\&undefined \ \xrightarrow{\ e\ } H,l,true$$

$$\frac{\text{Type(va1)} = \text{Number} \quad \text{Type(va2)} = \text{String}}{\text{H,l,va1==va2} \xrightarrow{e} \text{H,l,va1==@TN(va2)}}$$

$$\frac{\text{Type(va1)} = \text{Null} \quad \text{Type(va2)} != \text{Undefined} \quad \text{Type(va1)} != \text{Type(va2)}}{\text{H,l,va1==va2} \xrightarrow{e} \text{H,l,false}}$$

2.4 Statements

Similarly to the case for expressions, the semantics of statements contains a certain number of internal statements, used to represent unobservable execution steps, and user statements (reported in Figure 3)that are part of the user syntax of JavaScript.

---

s ::=
    "{"s*"}" *% block*
    var [(x["="e])˜] *% assignment*
    ; *% skip*
    e *% expression not starting with "",""function"*
    if "("e")" s [else s] *% conditional*
    while "("e")" s *% while*
    do s while "("e")"; *% do-while*
    for "("e in e")" s *% for-in*
    for "("var x["="e] in e")" s *% for-var-in*
    continue [x]; *% continue*
    break [x]; *% break*
    return [e]; *% return*
    with "("e")" s *% with*
    id:s *% label*
    throw e; *% throw*
    try "{"s*"}" [catch "("x"){"s1*"}"] [finally "{"s2*"}"] *% try*

Fig. 3. Syntax for Statements

A completion is the final result of evaluating a statement.

co ::= "("ct,vae,xe")" vae ::= &empty | va xe ::= &empty | x
ct ::= Normal | Break | Continue | Return | Throw

The completion type indicates whether the execution flow should continue normally, or be disrupted. The value of a completion is relevant when the completion type is Return (denoting the value to be returned), Throw (denoting the exception thrown), or Normal (propagating the value to be return during the execution of a function body). The identifier of a completion is relevant when the completion type is either Break or Continue, denoting the program point where the execution flow should be diverted to.

Expression and Throw. Evaluation contexts transform the expression operand of these constructs into a pure value. All the semantic rules specify is how such value is packaged into a completion. The rules for return,continue,break are similar.

$$\text{H,l,va} \xrightarrow{s} \text{H,l,(Normal,va,\&empty)} \qquad \text{H,l,throw va;} \xrightarrow{s} \text{H,l,(Throw,va,\&empty)}$$

Block Statements. The semantics of blocks handles the completions generated by the previous statements. A block does not introduce a new scope, as can be seen by the rules below. Rule (5) initializes the internal @Block statement so that it can always assume to have at least a statement inside (it uses the completion (Normal,&empty,&empty) which is equivalent to the empty statement ";"). Note that @Block(co,sC s∗) is a statement contexts sC, which evaluates nested statements. Rule (6) returns the result of evaluating a block to the top level, rule (7) leaves the block when a statement evaluates to an abrupt completion and rule (8) saves the last non-empty value and executes the next statement.

$$H,I,\{s*\} \xrightarrow{s} H,I,@Block((Normal,\&empty,\&empty)[,s+]) \quad (5)$$

$$H,I,@Block((ct,vae,xe)) \xrightarrow{s} H,I,(ct,vae,xe) \quad (6)$$

$$\frac{ct \ !< \ \{Normal\}}{H,I,@Block(co,(ct,vae,xe) \ s*) \xrightarrow{s} H,I,(ct,vae,xe)} \quad (7)$$

$$\frac{vae = IF \ vae2=\&empty \ THEN \ vae1 \ ELSE \ vae2}{H,I,@Block((ct,vae1,xe),(Normal,vae2,xe1) \ s*) \xrightarrow{s} H,I,@Block((Normal,vae,xe1)[,s+])} \quad (8)$$

With. The with statement is notoriously controversial. Its semantics is easy to explain on a technical level, but it is hard to have the correct intuition when writing code.

$$\frac{SetScope(I,I1,H)=H1,ln1,ln2}{H,I,with \ (I1) \ s \ \longrightarrow H1,I1,@with(I,ln1,ln2,s)}$$

$$H,I1,@with(I,null,null,co) \ \longrightarrow H,I,co$$

$$\frac{H(I1.@Scope=I2\{\})=H1}{H,I1,@with(I,I2,null,co) \ \longrightarrow H1,I,co}$$

$$\frac{H(I1.@Scope=I2\{\})=H1 \qquad H1(I3.@Scope=I1\{\})=H2}{H,I1,@with(I,I2,I3,co) \ \longrightarrow H2,I,co}$$

To evaluate the with construct we put its object parameter I1 as the top of the scope chain, and we save the current scope I so that it can be restored on exit. Moreover, we make use of the auxiliary function SetScope to deal with the case where I1 was already part of an existing scope chain.

This construct was introduced to give a convenient syntactic notation to operate on the properties of its object parameter. In practice though the statement with (o){ x = e } corresponds to if (o.hasOwnProperty("x")) x=e; else x=e, meaning that an attempt to update a non-existing property of o creates a global variable with the same name, rather than adding a property to o. The fault is not entirely on the with side though. For example the Byzantine semantics of the variable statement is responsible for other surprising behavior.

```
js> var n = {m:0}; with (n) {var m = 1}; n.m  % res: 1
js> m === undefined  % res: true
```

Above, var m = 1 is parsed before executing the program. It initializes m to &undefined in the global scope. At run time, the statement var m = 1 is essentially equivalent to the statement m = 1, that sets the property m of n to 1. The only difference is that var m=1 evaluates to a completion with value &empty, whereas m = 1 to one with value 1.

Try-catch. JavaScript provides a try-catch-finally mechanism to handle native and user-generated exceptions. We focus on the rule that does the interesting work, and that may raise some eyebrows.

$$\frac{\text{o=new\_object}(\textit{"Object"},\#\text{ObjectProt}) \quad \text{H1,l1=alloc(H,o)}}{\text{H2=H1(l1.}\textit{"x"}\text{=va\{DontDelete\})} \quad \text{H3=H2(l1.@Scope=l\{\})}}$$
$$\overline{\text{H,l,try (Throw,va,xe) catch (x) \{s1\}} \xrightarrow{s} \text{H3,l1,@catch \{s1\}}}$$

If the evaluation of the "tried" statement results in an exception, the catch code is executed in a new scope containing a binding of x to the exception value va. Combined with the semantics of methods by self-application, comes the surprising ability of JavaScript programs of tamper with their own scope:

```
js> x=0; function spy(){return this};
js> try {throw spy} catch(spy){spy().x = 1; x === 1} % res: true
js> x % res: 0
```

This snapshot is from the Rhino shell (see Section 3), which respect the ECMA specification. Implementations diverge on this point.

2.5 Programs

Programs are sequences of statements and function declarations.

$$\text{P ::= fd [P] | s [P]} \qquad\qquad \text{fd ::= function x } \textit{"("}[\text{x}\tilde{~}]\textit{")"}\{\textit{"[P]"}\}\textit{"}$$

As usual, the execution of statements is taken care of by a contextual rule. If a statement evaluates to a break or continue outside of a control construct, an SyntaxError exception is thrown (rule (9)). The run-time semantics of a function declaration instead is equivalent to a no-op (rule (10)). Function (and variable) declarations should in fact be parsed once and for all, before starting to execute the program text. In the case of the main body of a JavaScript program, the parsing is triggered by rule (11) which adds to the initial heap NativeEnv first the variable and then the function declarations (functions VD,FD).

$$\frac{\begin{array}{c}\text{ct} < \{\text{Break,Continue}\}\\ \text{o = new\_SyntaxError()} \quad \text{H1,l1 = alloc(H,o)}\end{array}}{\text{H,l,(ct,vae,xe) [P]} \xrightarrow{P} \text{H1,l,(Throw,l1,\&empty)}} \quad (9)$$

$$\text{H,l,function x ([x}\tilde{~}]\text{)\{[P]\} [P1]} \xrightarrow{P} \text{H,l,(Normal,\&empty,\&empty) [P1]} \quad (10)$$

$$\frac{\begin{array}{c}\text{VD(NativeEnv,\#Global,\{DontDelete\},P) = H1}\\ \text{FD(H1,\#Global,\{DontDelete\},P) = H2}\end{array}}{\text{P} \xrightarrow{P} \text{H2,\#Global,P}} \quad (11)$$

2.6 Native Objects

The initial heap NativeEnv of core JavaScript contains native objects for representing predefined functions, constructors and prototypes, and the global object @Global that constitutes the initial scope, and is always the root of the scope chain. As an example,

we describe the global object. The global object defines properties to store special values such as &NaN, &undefined etc., functions such as eval, toString etc. and constructors that can be used to build generic objects, functions, numbers, booleans and arrays. Since it is the root of the scope chain, its @Scope property points to null. Its @this property points to itself: @Global = {@Scope:null, @this:#Global, "eval":#GEval{DontEnum},...}. None of the non-internal properties are read-only or enumerable, and most of them can be deleted. By contrast, when a user variable or function is defined in the top level scope (i.e. the global object) it has only the DontDelete attribute. The lack of a ReadOnly attribute on "NaN","Number" for example forces programmers to use the expression 0/0 to denote the real &NaN value, even though @Number.NaN stores &NaN and is a read only property.

Eval. The eval function takes a string and tries to parse it as a legal program text. If it fails, it throws a SyntaxError exception (rule (12)). If it succeeds, it parses the code for variable and function declarations (respectively VD,FD) and spawns the internal statement @cEval (rule (13)). In turn, @cEval is an execution context for programs, that returns the value computed by the last statement in P, or &undefined if it is empty.

$$\frac{\text{ParseProg(m)} = \&\text{undefined} \qquad \text{H2,l2} = \text{alloc(H,o)} \quad o = \text{new\_SyntaxError()}}{\text{H,l,\#GEval.@Exe(l1,m)} \xrightarrow{e} \text{H2,l,<l2>}} \quad (12)$$

$$\frac{\begin{array}{c} l \mathrel{!=} \#\text{Global} \quad \text{ParseProg(m)} = P \\ \text{VD(H,l,\{\},P)} = \text{H1} \quad \text{FD(H1,l,\{\},P)} = \text{H2} \end{array}}{\text{H,l,\#GEval.@Exe(l1,m)} \xrightarrow{e} \text{H2,l,@cEval(P)}} \quad (13)$$

$$\frac{\text{va} = (\textit{IF}\ \text{vae=\&empty}\ \textit{THEN}\ \&\text{undefined}\ \textit{ELSE}\ \text{vae})}{\text{H,l,@cEval((ct,vae,xe))} \xrightarrow{e} \text{H,l,va}} \quad (14)$$

As we are not interested in modeling the parsing phase, we just assume a parsing function ParseProg(m) which given string m returns a valid program P or else &undefined. Note how in rule (13) the program P is executed in the caller's scope l, effectively giving dynamic scope to P.

Object. The @Object constructor is used for creating new user objects and internally by constructs such as object literals. Its prototype @ObjectProt becomes the prototype of any object constructed in this way, so its properties are inherited by most JavaScript objects. Invoked as a function or as a constructor, @Object returns its argument if it is an object, a new empty object if its argument is undefined or not supplied, or converts its argument to an object if it is a string, a number or a boolean. If the argument is a host object (such as a DOM object) the behavior is implementation dependent.

$$\frac{o = \text{new\_object}(\textit{"Object"},\#\text{ObjectProt}) \qquad \text{H1,lo} = \text{Alloc(H,o)} \quad \text{Type(pv)} < \{\text{Null,Undefined}\}}{\text{H,l,\#Object.@Construct(pv)} \xrightarrow{e} \text{H1,l,lo}}$$

$$\frac{\text{Type(pv)} < \{\text{String, Boolean, Number}\} \quad \text{H1,le} = \text{ToObject(H,pv)}}{\text{H,l,\#Object.@Construct(pv)} \xrightarrow{e} \text{H,l,le}} \qquad \frac{\text{Type(l1)} = \text{Object} \quad !\text{IsHost(H,l1)}}{\text{H,l,\#Object.@Construct(l1)} \xrightarrow{e} \text{H,l,l1}}$$

The object @ObjectProt is the root of the scope prototype chain. For that reason, its internal prototype is null. Apart from "constructor", which stores a pointer to @Object,

the other public properties are native meta-functions such as toString or valueOf (which, like user function, always receive a value for @this as the first parameter).

Native functions are always built using the macro

@fun,@funProt = make_native_fun(#fun,#funProt,n)

which creates a skeleton object for @fun and its prototype at the addresses #fun,#funProt, and declares n to be the expected number of actual parameters. For example,

@OPtoString,@OPtoStringProt = make_native_fun(#OPtoString,#OPtoStringProt,0)

$$\frac{H(l1).@Class = m}{H,l,\#OPtoString.@Exe(l1) \xrightarrow{e} H,l,"[object"m"]"}$$

@OPvalueOf,@OPvalueOfProt = make_native_fun(#OPvalueOf,#OPvalueOfProt,0)

$$\frac{!IsHost(H,l1)}{H,l,\#OPvalueOf.@Exe(l1) \xrightarrow{e} H,l,l1}$$

@OPtoString returns a rather uninformative string showing the class (i.e. essentially the constructor) of its argument, and @OPvalueOf returns directly its argument.

js> var t = {l:0}; t.valueOf() === t % res: true
js> t % res: [object Object]

Function. The Function constructor is rarely used, as functions are usually created by function declarations or function expression, but its prototype is the internal prototype of all functions, no matter how they are created.

The Function constructor tries to parse its list of parameters into a (possibly empty) comma-separated list of formal parameters, and a program text. If the parsing succeeds, it returns a pointer to a newly allocated function.

$$H,l,\#Function.@Construct([va\tilde{\ },]va) \xrightarrow{e} H,l,@FunParse("",[va\tilde{\ },]va)$$

$$\frac{\begin{array}{l} ParseParams("x\tilde{\ }") \\ ParseFunction("P") \\ H,Function(fun(x\tilde{\ })P,\#Global) = H1,l1 \end{array}}{H,l,@FunParse("x\tilde{\ }","P") \xrightarrow{e} H1,l,l1}$$

Oddly, instead of its proper lexical scope, the function gets object @Global as its scope, by the expression H,Function(fun(x˜){P},#Global) = H1,l1). Hence,

js> (function (){var b=1; var c= function (){x=b}; c()})(); x % res: 1
js> (function (){var b=1; var c= new Function("x=b"); c()})() % res: ReferenceError: b is not defined

The Function object prototype @FunctionProt is a function itself, which just returns @undefined and defines fields such as

@FunctionProt = {
    @Class: "Function",
    @Prototype: #ObjectProt,
    "constructor": #Function{DontEnum},
    "prototype": #FunctionProtProt{DontDelete},

```
    "call": #FPcall{DontEnum}
    ...}
```

$$H,I,\#FunctionProt.@Exe(I1) \xrightarrow{e} H,I,@undefined$$

Its method *"call"* its very characteristic of JavaScript, exposing the very essence of the language to the programmer. JavaScript in fact is, at the core, a functional language where objects are records, and where functions are first-class values. The typical behavior of object-oriented languages is achieved by parameterizing functions by an implicit first argument which represents the object/record of which they are a method. The method invocation o.f(v) is essentially equivalent to (o.f).call(o,v), where first we extract function f from o, and then we use the call method to apply f to o itself, and to its explicit parameter v. The rules below illustrate the behavior of the call method.

$$\frac{\begin{array}{c} \mathsf{GetType}(H,I1) = \mathit{"function"} \\ \mathsf{va} < \{\mathsf{null},\#\mathsf{undefined}\} \end{array}}{H,I,\#\mathsf{FPcall}.@\mathsf{Exe}(I1,\mathsf{va}[,\mathsf{va}^\sim]) \xrightarrow{e} H,I,I1.@\mathsf{Call}(\#\mathsf{Global}[,\mathsf{va}^\sim])}$$

$$\frac{\begin{array}{c} \mathsf{GetType}(H,I1) = \mathit{"function"} \\ \mathsf{va} \:!< \{\mathsf{null},\#\mathsf{undefined}\} \end{array}}{H,I,\#\mathsf{FPcall}.@\mathsf{Exe}(I1,\mathsf{va}[,\mathsf{va}^\sim]) \xrightarrow{e} H,I,I1.@\mathsf{Call}(\mathsf{va}[,\mathsf{va}^\sim])}$$

By de-sugaring method invocation, #FPCall makes it possible to turn arbitrary functions into methods without modifying the target object. For example, we can use the standard *"toString"* method of object on arrays, which have their own different *"toString"*.

```
js> var a = [1,2,3]; a.toString() % res: 1,2,3
js> toString.call(a) % res: [object Array]
```

3 Relation to Implementations

So far we have referred to JavaScript as the "abstract" language defined by the ECMA specification. In practice, JavaScript is implemented within the major web browsers, or as standalone shells. In order to clarify several ambiguities present in the specification, we have run experiments inspired by our semantics rules on different implementations. We have found that, besides resolving ambiguities in different and often incompatible ways, implementations sometimes do not conform to the specification in implementing corner (and not so corner) cases which have a well-defined semantics.

When the specification states that some behavior is implementation dependent, we either choose one, or parameterize the semantics by some partially unspecified helper function.

Moreover, Mozilla maintains and develops the original JavaScript language (embodied in its C and Java implementations, Spidermonkey and Rhino) which extends the functionalities of ECMA-262 in several ways. In this section we describe, mainly by examples, some interesting differences and incompatibilities between JavaScript implementations. We conclude with a brief discussion on how some of the JavaScript 1.5 extensions may be integrated with our framework.

3.1 Divergence in implementations

Below, the prompt "?>" is used in examples to denote an hypothetical strictly ECMA-262 compliant implementation.

Function expressions and declarations. This is one of the most striking examples of implementations diverging from the specification. As prescribed by the specification, an expression cannot be parsed as a statement if it starts with function, hence s1; function a(){}; s2 can be parsed as program but cannot be parsed as part of a block statement. Function definitions, according to the specification, are parsed in source text order before running the program code, and are added to the scope. Hence, the code below is fully specification compliant in choosing the second definition in the first examples, and throwing an exception in the second.

```
?> (function (){function b(){return 1}; return b();
                function b(){return 0}})()  % res: 0
?> (function (){{function b(){return 1}}; return b()})()  % res: SyntaxError: ...
```

Major implementations instead accept function definitions as statements, with mixed results. From our experiments, we think that Rhino parses standard-compliant function definitions according to the specification, removes them from the parsed code, but considers at run time (with the same semantics) the remaining function definitions that appear as statements.

```
js> (function (){function b(){return 0};
          if (true){function b(){return 1}}; return b()})()  % res: 1
js> (function (){function b(){return 0};
          if (false){function b(){return 1}}; return b()})()  % res: 0
```

We find this behavior completely reasonable, and easy to implement in a formal semantics.

Our experiments on SpiderMonkey gave different results. Apparently, the semantics of functions in this implementation depends on the position, *within unreachable code* of statements (such as var g below) which should have no semantic significance!

```
js> (function(){if (true) function g(){return 0};
                  return g(); var g; function g(){return 1}})()  % res: 0
js> (function(){if (true) function g(){return 0};
                  return g(); function g(){return 1}; var g})()  % res: 1
```

Moreover, SpiderMonkey mishandles the attributes of function declarations interpreted as statements in eval code, where functions and variables should be deleted at will.

```
js> eval("function a(){};delete a")  % res: true
js> eval("function a(){};if(1){function a(){return 1}};delete a")  % res: false
```

JScript parses all the functions declarations at the beginning, including the ones appearing as statements.

```
> if(true){function a(){return 1}}else{function a(){return 0}};a()  % res: 0
```

Named function expressions. Function declarations bind a variable in the global scope, so the correct way to implement recursive functions is by using named function expressions, whose names have local scope within the function.

```
js> function a(){a()}; c=a; a = 1; c()  % res: TypeError: a is not a function
js> var a = function b(){b()}; c=a; a=1; c()  % res: InternalError: too much recursion
js> b  % res: ReferenceError: b is not defined
```

JScript unfortunately binds also *b* in the global environment, so assigning to *b* we can disrupt the recursive definition.

Internal creation of new objects. We have shown in Section 2.3 that literal objects are created by evaluating the expression new Object(), instead of creating internally a new empty object *as described in the specification for* new Object. The specification is ambiguous regarding this and several similar cases. We followed the main implementations in choosing where to evaluate the expression new Object() and when to use a fresh empty object as described in the specification of new Object. For example, the latter option is popular for creating the scope object in the catch branch of a try-catch construct.

Unfolding of while loops. The specification prescribes that the top level statements in a program text do not propagate their computed value. Hence, a hypothetical implementation would not respect for example the simple minded equivalence on while unfolding shown below, because the eval code in the second example (which must be considered as program text) terminates with a while loop returning the internal &empty value, that as we have seen in Section 2.6 is converted to &undefined.

```
?> eval("x=1;while␣(x−−␣>0){y=1}") === 1  % res: true
?> eval("x=1;if(x−−␣>0)y=1;while␣(x−−␣>0){y=1}") === undefined  % res: true
```

Functions returning references. The specification makes it clear that certain host functions may return values of a reference type. Apparently there exist some function in the Internet Explorer implementation that does so (we were not able to find which one). Our semantics is compatible with this behavior, but we have found that this corner case complicates the formalization in several points, in particular by not granting the assumption that values returned by expressions are in fact *pure values*, and hence having to constantly refere to the pervasive GetValue function.

Expressions as statements. The specification states that an expression can be regarded as a statement as long as it does not start with either "{" or function. An unfortunate consequence is that the naive programmer may think that he is writing an object literal where he is in face defining a new global function, by using a labelled statement inside a block.

```
js> {l:function f(){}}; f  % res: function f()
```

Native functions used as constructors. The specification dictates that native functions cannot be used as constructors unless explicitly mentioned in their description. We reflected this in our semantics by not providing native functions (such as eval) with a @Construct method, but many implementations have chosen not to do so.

```
js> var a = new eval()
js> a  % res: [object Object]
js> var a = new Function.prototype()
js> a  % res: [object Object]
```

Function arguments aliasing. The actual parameters of a function call are collected in an arguments array-like object available inside the function scope. The fields of this array are aliases for the formal parameters:

```
js> (function (x){arguments[0]=1; return x})(4) % res: 1
```

This is an odd feature, the only case where such kind of aliasing is observable in JavaScript. This feature is going to be removed form future versions of the standard, and we have decided not to model it in our semantics because it requires extra machinery, and does not seem to affect the semantics of programs in an interesting way, apart perhaps from confusing the programmer. Standard implementations seem to be using getters and setters to model this property.

```
js> (function(x){delete arguments[0];arguments[0]=1;return x})(4) % res: 4
```

We do not believe that this behavior is entirely justified, as the specification can be interpreted as requiring that the aliasing should be in place whenever arguments has a property corresponding to the position of an existing formal parameter.

Joined objects. The specification provides for the possibility that functions "defined by the same piece of source text" be implemented as joined objects, i.e. sharing their properties. If that were the case, we could have

```
?> function f(){function g(){}; return g}
?> var h = f(); var j = f(); h.a = 0; i.a % res: 0
```

Fortunately, no known implementation uses this "feature", and we could scrap it from our semantics. The possibility to have joined objects is going to be removed from future versions of the language.

Scoping of the catch construct. We have shown in Section 2.4 that the scoping mechanism of the try-catch construct can lead to programs getting hold of their own scope. SpiderMonkey and other implementations decided to protect the programmer from such abomination, and pass instead the global object as the implicit this parameter of the spy function below.

```
js> x=0; function spy(){return this};
js> try {throw spy} catch(spy){spy().x = 1; x === 1} % res: true
js> x % res: 1
```

3.2 JavaScript1.5 Extensions

Each major web browser and shell has its own custom extensions to JavaScript. We discuss briefly the case of getters and setters (introduced in Mozilla's JavaScript 1.5) and related properties, because they are the most interesting from a semantics viewpoint, and have the greatest impact in terms of making JavaScript harder to analyze.

A (property) *getter* is essentially a function that gets called when the corresponding property is accessed, and a *setter* is a function that gets called when the property is assigned. User-defined getters and setters allow a programmer to override the normal property access and assignment functions. For example:

```
js> var o = {count:0,get p(){this.count=this.count+1; return 0},};
js> o.count % res: 0
```

js> `o.p` *% res: 0*
js> `o.count` *% res: 1*

We propose to model getters and setters for a property p by creating a dummy object with internal properties `@Getter` and `@Setter` pointing to the corresponding function objects. Functions `@Get` and `@Put` need to change so that when the property is accessed, the function first checks if the property has getters and setters (which can be done by checking if the property points to an object containing the `@Getter` and `@Setter` properties) and then calls the appropriate functions. We leave the full integration of this approach into our operational semantics to future work.

4 Formal Properties

In this section we give some preliminary definitions and set up a basic framework for formal analysis of well-formed JavaScript programs. We prove a progress theorem which shows that the semantics is sound and the execution of a well-formed term always progresses to an exception or an expected value. Next we prove a Heap reachability theorem which essentially justifies mark and sweep type garbage collection for JavaScript. Although the properties we prove are fairly standard for idealized languages used in formal studies, proving them for real (and unruly!) JavaScript is a much harder task.

Throughout this section, a program state S denotes a triple (H, l, t) where H is a heap, l is a current scope address and t is a term. Recall that a heap is a map from heap addresses to objects, and an object is a collection of properties that can contain heap addresses or primitive values.

4.1 Notation and Definitions

Expr, *Stmnt* and *Prog* denote respectively the sets of all possible expressions, statements and programs that can be written using the corresponding internal and user grammars. $\mathbb{L}$ denotes the set of all possible heap addresses. *Wf(S)* is a predicate denoting that a state S is well-formed. *prop(o)* is the set of property names present in object o. *dom(H)* gives the set of allocated addresses for the heap H.

For a heap address l and a term t, we say $l \in t$ iff the heap address l occurs in t. For a state $S = (H, l, t)$, we define $\Delta(S)$ as the set of heap addresses $\{l\} \cup \{l | l \in t\}$. This is also called the set of *roots* for the state S.

We define the well-formedness predicate of a state $S = (H, l, t)$ as the conjunction of the predicates $Wf_{Heap}(H)$, $Wf_{scope}(l)$ and $Wf_{term}(t)$. A term t is well-formed iff it can be derived using the grammar rules consisting of both the language constructs and the internal constructs, and all heap addresses contained in t are allocated ie $l \in t \Rightarrow l \in dom(H)$. A scope address $l \in dom(H)$ is well-formed iff the scope chain starting from l does not contain cycles, and $(@Scope \in prop(H(l))) \wedge (H(l).@Scope \neq null \Rightarrow \mathrm{Wf}(H(l).@Scope))$. A heap H is well-formed iff it conforms to all the conditions on heap objects mentioned in the specification (see Section A for a detailed list).

Definition 1 (Heap Reachability Graph). *Given a heap H, we define a labeled directed graph G_H with heap addresses $l \in dom(H)$ as the nodes, and an edge from address l_i to l_j with label p iff $(p \in prop(H(l_i)) \wedge H(l_i).p = l_j)$.*

Given a heap reachability graph G_H, we can define the view from a heap address l as the subgraph $G_{H,l}$ consisting only of nodes that are reachable from l in graph G_H. We use $\text{view}_H(l)$ to denote the set of heap addresses reachable from l: $\text{view}_H(l) = Nodes(G_{H,l})$. view_H can be naturally extended to apply to a set of heap addresses. Observe that the graph G_H only captures those object properties that point to other heap objects and does not say anything about properties containing primitive values.

Definition 2 (l-Congruence of Heaps $\cong_l$). *We say that two heaps H_1 and H_2 are l-congruent (or congruent with respect to heap address l) iff they have the same* views *from heap address l and the corresponding objects at the heap addresses present in the views are also equal. Formally,*

$$H_1 \cong_l H_2 \Leftrightarrow (G_{H_1,l} = G_{H_2,l} \ \wedge \ \forall \ l' \in \ \text{view}_{H_1}(l) \ H_1(l') = H_2(l')).$$

Note that if $H_1 \cong_l H_2$ then $\text{view}_{H_1}(l) = \text{view}_{H_2}(l)$. It is easy to see that if two heaps H_1 and H_2 are congruent with respect to l then they are congruent with respect to all heap addresses $l' \in \text{view}_{H_1}(l)$.

Definition 3 (State congruence $\cong$). *We say that two states $S_1 = (H_1, l, t)$ and $S_2 = (H_2, l, t)$ are congruent iff the heaps are congruent with respect to all addresses in the roots set. Formally, $S_1 \cong S_2 \Leftrightarrow \forall \ l' \in \Delta(S_1) \ (H_1 \cong_{l'} H_2)).$*

Note that $\Delta(S_1) = \Delta(S_2)$ because the definition of Δ depends only on l and t. In the next section we will show that for a state $S = (H, l, t)$, $\text{view}_H(\Delta(S))$ forms the set of *live* heap addresses for the S because these are the only possible heap addresses that can be accessed during any transition from S.

Definition 4 (Heap Address Renaming). *For a given heap H, a heap address renaming function f is any one to one map from $\text{dom}(H)$ to $\mathbb{L}$.*

We denote the set of all possible heap renaming functions for a heap H by $\mathbb{F}_H$. We overload f so that $f(H)$ is the new heap obtained by renaming all heap addresses $l \in \text{dom}(H)$ by $f(l)$ and for a term t, $f(t)$ is the new term obtained by renaming all $l \in t$ by $f(l)$. Finally, for a state $S = (H, l, t)$ we define $f(S) = (f(H), f(l), f(t))$ as the new state obtained under the renaming.

Definition 5 (State similarity $\sim$). *Two states $S_1 = (H_1, l_1, t_1)$ and $S_2 = (H_2, l_2, t_2)$ are similar iff there exists a renaming function f for H_1 such that the new state $f(S_1)$ obtained under the renaming is congruent to S_2. Formally,*

$$S_1 \sim S_2 \Leftrightarrow \exists \ f \in \mathbb{F}_{H_1} \ f(S_1) \cong S_2.$$

Property 1. Both $\cong$ and $\sim$ are equivalence relations. Moreover, $\cong \subsetneq \sim$.

4.2   Theorems and Formal Properties

We now present the main technical results. Our first result is a progress and preservation theorem, showing that evaluation of a well-formed term progresses to a value or an exception.

Lemma 1. *Let C denote the set of all valid contexts for expressions, statements and programs. For all terms t appropriate for the context C we have $Wf_{term}(C(t)) \Rightarrow Wf_{term}(t)$.*

Theorem 1 (Progress and Preservation). *For all states $S = (H, l, t)$ and $S' = (H', l', t')$:*

- *$(Wf(S) \wedge S \to S') \Rightarrow Wf(S')$ (Preservation)*
- *$Wf(S) \wedge t \notin v(t) \Rightarrow \exists S' (S \to S'))$ (Progress)*

where $v(t) = ve$ if $t \in Expr$ and $v(t) = co$ if $t \in Stmnt$ or $Prog$.

Our second result shows that similarity is preserved under reduction, which directly gives a construction for a simple mark-and-sweep-like garbage collector for JavaScript. The proofs for the theorems are given in Section A.

Lemma 2. *For all well-formed program states $S = (H, l, t)$, if $H, l, t \to H', l', t'$ then $H(l'') = H'(l'')$, for all $l'' \notin view_H(\Delta(H, l, t)) \cup view_{H'}(\Delta(H', l', t'))$.*

The above lemma formalizes the fact that the only heap addresses accessed during a reduction step are the ones present in the initial and final live address sets. We can formally prove this lemma by an induction over the rules.

Theorem 2 (Similarity preserved under reduction). *For all well-formed program states S_1, S_2 and S_1', $S_1 \sim S_2 \wedge S_1 \to S_1' \Rightarrow \exists S_2'. S_2 \to S_2' \wedge S_1' \sim S_2'$.*

We can intuitively understand this theorem by observing that if the reduction of a term does not involve allocation of any new heap addresses then the only addresses that can potentially be accessed during the reduction would be the ones present in the live heap address set. When the program states are similar, then under a certain renaming the two states would have the same live heap address sets . As a result the states obtained after reduction would also be congruent(under the same renaming function). On the other hand, if the reduction involves allocation of new heap addresses then we can simply extend the heap address renaming function by creating a map from the newly allocated addresses in the first heap (H_1') to the newly allocated addresses in the second heap (H_2'). Thus state similarity would be preserved in both cases.

A consequence of Theorem 2 is that we can build a simple mark and sweep type garbage collector for JavaScript. For any program state $S = (H, l, t)$, we mark all the heap addresses that are reachable from $\Delta(S)$. We modify the heap H to H' by freeing up all unmarked addresses and obtain the new program state $S' = (H', l, t)$. It is easy to show that $S' \sim S$. Hence by Theorem 2, a reduction trace starting from t, in a system with garbage collection, would be similar to the one obtained without garbage collection. In other words, garbage collection does not affect the semantics of programs.

5   Related Work

The JavaScript approach to objects is based on Self [26] and departs from the foundational object calculi proposed in the 1990s, e.g., [5,9,18]. Previous foundational studies include operational semantics for a subset of JavaScript [12] and formal properties of subsets of JavaScript [7,22,25,24]. Our aim is different from these previous efforts because we address the full ECMA Standard language (with provisions for variants introduced in different browsers). We believe a comprehensive treatment is important for analyzing existing code and code transformation methods [1,2]. In addition, when analyzing JavaScript security, it is important to consider attacks that could be created

using arbitrary JavaScript, as opposed to some subset used to develop the trusted application. Some work on containing the effects of malicious JavaScript include [21,27]. Future versions of JavaScript and ECMAScript are documented in [14,13].

In the remainder of this section, we compare our work to the formalizations proposed by Thiemann [25] and Giannini [7] and comment on the extent to which formal properties they establish for subsets of JavaScript can be generalized to the full language.

Giannini et al. [7] formalize a small subset of JavaScript and give a static type system that prevents run-time typing errors. The subset is non-trivial, as it includes dynamic addition of properties to objects, and constructor functions to create objects. However the subset also lacks important features such as object prototyping, functions as objects, statements such as with, try−catch, for−in, and native functions and objects. This leads to substantial simplifications in their semantics, relative to ours. For example, function definitions are stored in a separate data structure rather than in the appropriate scope object, so there is no scope-chain-based resolution of global variables appearing inside a function body. Their simplification also makes it possible to define a sound type system that does not appear to extend to full JavaScript, as further discussed below.

Thiemann [25] proposes a type system for a larger subset of JavaScript than [7], as it also includes function expressions, function objects, and object literals. The type system associates type signatures with objects and functions and identifies suspicious type conversions. However, Thiemann's subset still does not alllow object prototyping, the with and the try−catch statements, or subtle features of the language such as property attributes or arbitrary variable declarations in the body of a function. As we showed in section 2, these non-trivial (and non-intuitive) aspects of JavaScript make static analysis of arbitrary JavaScript code very difficult.

The substantial semantic difference between the subsets covered in [7,25] and full JavaScript is illustrated by the fact that: (i) Programs that are well-typed in the proposed subsets may lead to type errors when executed using the complete semantics, and (ii) Programs that do not lead to a type error when executed using the complete semantics may be considered ill-typed unnecessarily by the proposed type systems. The first point is demonstrated by var x = "a"; x.length = function(){ }; x.length(), which is allowed and well-typed by [25]. However it leads to a type error when executed using the complete semantics because although the type system considers implicit type conversion of the string x to a wrapped string object, it does not consider the prototyping mechanism and attributes for properties. Since property length of String.prototype has the ReadOnly attribute, the assignment in the second statement fails silently and thus the method call in the third statement leads to a type error. An example demonstrating the second point above is function f(){return o.g();}; result = f(), which is allowed in the subsets mentioned in [7,25]. If method g is not present in the object o then both type systems consider the expression result = f() ill-typed. However g could be present as a method in the one of the ancestoral prototypes of o, in which case the expression will not lead to a type error. Because object prototyping is the main inheritance mechanism in JavaScript and it is pervasive in almost all real world JavaScript, we believe that a type system that does not consider the effects of prototypes will not be useful without further extension.

6 Conclusions

In this paper, we describe a structured operational semantics for the ECMA-262 standard [15] language. The semantics has two main parts: one-step evaluation relations for the three main syntactic categories of the language, and definitions for all of the native objects that are provided by an implementation. In the process of developing the semantics, we examined a number of perplexing (to us) JavaScript program situations and experimented with a number of implementations. To ensure accuracy of our semantics, we structured many clauses after the ECMA standard [15]. In a revision of our semantics, it would be possible to depart from the structure of the informal ECMA standard and make the semantics more concise, using many possible optimization to reduce its apparent complexity. As a validation of the semantics we proved a soundness theorem, a characterization of the reachable portion of the heap, and some equivalences between JavaScript programs.

In ongoing work, we are using this JavaScript semantics to analyze methods for determining isolation between embedded third-party JavaScript, such as embedded advertisements provided to web publishers through advertising networks, and the hosting content. In particular, we are studying YAHOO!'s ADsafe proposal [1] for safe online advertisements, the BeamAuth [6] authentication bookmarklet that relies on isolation between JavaScript on a page and JavaScript contained in a browser bookmark,Google's Caja effort [2] to provide code isloation by JavaScript rewriting, and FBJS [23], the subset of JavaScript used for writing FaceBook applications.

References

1. ADsafe: Making JavaScript safe for advertising. `http://www.adsafe.org/`, May 2008.
2. Google-Caja: A source-to-source translator for securing JavaScript-based web. `http://code.google.com/p/google-caja/`.
3. JScrip (Windows Script Technologies).
   `http://msdn2.microsoft.com/en-us/library/hbxc2t98.aspx`.
4. Rhino: JavaScript for Java. `http://www.mozilla.org/rhino/`.
5. M. Abadi and L. Cardelli. *A Theory of Objects*. Springer, 1996.
6. B. Adida. Beamauth: two-factor web authentication with a bookmark. In *ACM Computer and Communications Security*, pages 48–57, 2007.
7. C. Anderson, P. Giannini, and S. Drossopoulou. Towards type inference for JavaScript. In *Proc. of COOP'05*, page 429452, 2005.
8. B. Eich. JavaScript at ten years. `www.mozilla.org/js/language/ICFP-Keynote.ppt`.
9. K. Fisher, F. Honsell, and J.C. Mitchell. A lambda calculus of objects and method specialization. *Nordic J. Computing (*formerly *BIT)*, 1:3–37, 1994.
10. D. Flanagan. *JavaScript: The Definitive Guide*. O'Reilly, 2006. `http://proquest.safaribooksonline.com/0596101996`.
11. P. Gardner, G. Smith, M. Wheelhouse, and U. Zarfaty. Local hoare reasoning about DOM. In *Proc. of PODS '08*, pages 261–270. ACM, 2008.
12. D. Herman. Classic JavaScript. `http://www.ccs.neu.edu/home/dherman/javascript/`.

13. D. Herman and C. Flanagan. Status report: specifying JavaScript with ML. In *Proc. of ML'07*, pages 47–52, 2007.

14. ECMA International. ECMAScript 4. `http://www.ecmascript.org`.

15. ECMA International. ECMAScript language specification. stardard ECMA-262, 3rd Edition.
`http://www.ecma-international.org/publications/files/ECMA-ST/Ecma-262.pdf`, 1999.

16. S. Maffeis, J. Mitchell, and A. Taly. Complete ECMA 262-3 operational semantics. `http://jssec.net/semantics/`.

17. S. Maffeis, J. Mitchell, and A. Taly. An operational semantics for javascript. In *Proc. of APLAS'08*, volume 5356 of *LNCS*, pages 307–325. Springer Verlag, December 2008.

18. J.C. Mitchell. Toward a typed foundation for method specialization and inheritance. In *Proc. of POPL'90*, pages 109–124, 1990.

19. Mozilla. SpiderMonkey (JavaScript-C) engine.
`http://www.mozilla.org/js/spidermonkey/`.

20. Prototype Core Team. Prototype Javascript framework: Easy Ajax and DOM manipulation for dynamic web applications. `http://www.prototypejs.org`.

21. C. Reis, J. Dunagan, H. Wang, O. Dubrovsky, and S. Esmeir. BrowserShield: Vulnerability-driven filtering of Dynamic HTML. *ACM Transactions on the Web*, 1(3), 2007.

22. J. Siek and W. Taha. Gradual typing for objects. In *ECOOP*, 2007.

23. The FaceBook Team. FBJS. `http://wiki.developers.facebook.com/index.php/FBJS`.

24. P. Thiemann. Towards a type system for analyzing javascript programs. In *Proc. of ESOP'05*, volume 3444 of *LNCS*, page 408422, 2005.

25. P. Thiemann. A type safe DOM API. In *Proc. of DBPL*, pages 169–183, 2005.

26. D. Ungar and R.B. Smith. Self: The power of simplicity. In *Proc. OOPSLA)*, volume 22, pages 227–242, 1987.

27. D. Yu, A. Chander, N. Islam, and I. Serikov. JavaScript instrumentation for browser security. In *Proc. of POPL'07*, pages 237–249, 2007.

A   Proofs from Section 4

Wellformedness of Heaps. We say that a heap H is well-formed iff the following conditions are true.

- Every object in the heap must have *@Class* and *@Prototype* in its set of properties.
- Every function object in the heap must have *@Call, @Scope, length, @Body* and *@Prototype* in its set of properties.
- Every arguments object in the heap must have *callee* and *length* in its set of properties.
- Every array object in the heap must have the *length* property. The length property must always contain a number.
- Every native function object must have the *@actuals* property.
- Every String, Number and Boolean object must have the *@Value* property.
- All native error objects must have the message property.
- #Global must be an allocated address and must atleast have the *@this* property.
- The prototype chain for any object must never contain a cycle.
- The scope property for any function object must contain a well-formed scope address and also the body property must contain a well-formed term.

Proof of Theorem 1.

Proof Sketch: We prove the first part of the above theorem by an induction over the rules. The base cases are the rules which do not have a state transition in their premise. For each of them, we show directly that well-formedness of the initial state implies well-formedness of the final state. The remaining (contextual) rules form our inductive cases, for which we argue inductively as follows. Consider the rule:

$$\frac{\mathsf{H,l,t} \longrightarrow \mathsf{H',l',t'}}{\mathsf{H,l,C(t)} \longrightarrow \mathsf{H',l',C(t')}}$$

Let $Sc = (H, l, C(t))$ and $S = (H, l, t)$. We can similarly define Sc' and S'. We know that $\mathrm{Wf}(Sc) = \mathrm{Wf}_{Heap}(H) \wedge \mathrm{Wf}_{scope}(l) \wedge \mathrm{Wf}_{term}(t)$. Using Lemma 1, we know that $\mathrm{Wf}_{term}(C(t)) \Rightarrow \mathrm{Wf}_{term}(t)$. Therefore $\mathrm{Wf}(Sc) \rightarrow \mathrm{Wf}(S)$. From the hypothesis we can conclude $\mathrm{Wf}(Sc) \rightarrow \mathrm{Wf}(S) \Rightarrow \mathrm{Wf}(S') \Rightarrow \mathrm{Wf}(H') \wedge \mathrm{Wf}(l') \wedge \mathrm{Wf}(t')$. Thus all that is left to show is that $\mathrm{Wf}_{term}(t') \Rightarrow \mathrm{Wf}_{term}(C(t'))$. To prove this we do a case analysis over all contextual terms $C(t)$ arising from the inductive cases, and show that if t is well-formed then $C(t)$ is also well-formed. Therefore $\mathrm{Wf}(Sc) \rightarrow \mathrm{Wf}(Sc')$.

The second part of the theorem can be proven by structural induction over the terms and some case analysis. All the base cases which are not values/exceptions, have a basic reduction rule (basic rules are those which don't have a state transition in their premise) that applies to them. So the theorem is true for the base cases. For the inductive case, we show that for each expression, statement and program either is a context rule that applies (the premise of the context rule would be true by the induction hypothesis), or the term is a value or an exception, in which case the theorem is directly true. $\square$

Proof of Theorem 2. In order to prove this theorem we prove the following stronger theorem : For program states $S_1 = (H_1, l_1, t_1)$ and $S_2 = (H_2, l_2, t_2)$,

$\exists f \in \mathbb{F}_{H_1} \ \exists f' \in \mathbb{F}_{H_1'} \ (S_1 \rightarrow S_1' \ \wedge \ f(S_1) \cong S_2) \Rightarrow (S_2 \rightarrow S_2' \ \wedge \ f'(S_1') \cong S_2' \ \wedge \ \forall l \in dom(H_1)) f(l) = f'(l) \ \wedge \ \forall \ (l \in (\Delta(S_2') \cup \Delta(S_2))) \ f'(H_1') \cong_l H_2')$

The above statement basically says that if the initial states are similar then after reduction the final states would be similar and the final heaps would also be congruent with respect to the initial set of roots ($\Delta(S)$) as well. We prove the above theorem by an induction over the depth of the derivation tree. All rules which do no have a state transition in their premise form the base case (as they lead to a single step derivation). For base cases we do a case analysis and show that the theorem is true. The contextual rules form the inductive cases for which we argue as follows. Reduction using a contextual rule would look like:

$$\frac{\text{H1,l1,t1} \longrightarrow \text{H1',l1',t1'}}{\text{H1,l1,C1(t1)} \longrightarrow \text{H1',l1',C1(t1')}} \qquad \frac{\text{H2,l2,t2} \longrightarrow \text{H2',l2',t2'}}{\text{H2,l2,C2(t2)} \longrightarrow \text{H2',l2',C2(t2')}}$$

Let $Sc_1 = (H_1, l_1, C_1(t_1))$, $S_1 = (H_1, l_1, t_1)$, $Sc_2 = (H_2, l_2, C_2(t_2))$, $S_2 = (H_2, l_2, t_2)$ and similarly define states Sc_1', $S1'$, Sc_2' and $S2'$. Let $\mathcal{L}_{C_1} = \{l \mid l \in C_1\}$ and $\mathcal{L}_{C_2} = \{l \mid l \in C_2\}$.

$Sc_1 \sim Sc_2 \Rightarrow \exists f \in \mathbb{F}_{H_1} f(Sc_1) \cong Sc_2$. This would imply

a. $\Delta(f(Sc_1)) = \Delta(Sc_2) \Rightarrow f(\mathcal{L}_{C_1}) = \mathcal{L}_{C_2}$
b. $\forall l \in \Delta(Sc_2)) f(H_1) \cong_l H_2$
c. $f(C_1(t_1)) = f(C_1)(f(t_1)) = C_2(t_2)$

From (c) we get, $f(t_1) = t_2 \wedge f(C_1) = C_2$. Combining this with (a) we get $\Delta(f(S_1)) = \Delta(S_2)$. Since $\Delta(S_2) \subseteq \Delta(Sc_2)$, using (b) we get $\forall l \in \Delta(S_2)) f(H_1) \cong_l H_2$ and hence $S_1 \sim S_2$. Using the induction hypothesis we have $\exists f' \in \mathbb{F}_{H_1'}$ such that

d. $\Delta(f'(S_1')) = \Delta(S_2')$
e. $f'(t_1') = t_2'$
f. $\forall l \in dom(H_1) \ f'(l) = f(l)$
g. $\forall l \in (\text{view}_{H_2}(\Delta(S_2)) \ \cup \ \text{view}_{H_2'}(\Delta(S_2')))f'(H_1') \cong_l H_2'$

Since $L_{C_1} \subseteq dom(H_1)$, using (f) we get $f(\mathcal{L}_{C_1}) = f'(\mathcal{L}_{C_1})$. Also it is easy to observe that :

h. $\Delta(f'(Sc_1')) = f'(\mathcal{L}_{C_1}) \ \cup \ \Delta(f'(S_1')) = f(\mathcal{L}_{C_1}) \cup \Delta(f'(S_1'))$
i. $\Delta(Sc_2') = \mathcal{L}_{C_2} \ \cup \ \Delta(S_2')$
j. $\Delta(Sc_2) = \mathcal{L}_{C_2} \ \cup \ \Delta(S_2)$

From (a), (d), (h) and (i) we get $\Delta(Sc_2') = \Delta(f'(Sc_1'))$. Combining this with (e) we get $f'(C_1(t_1')) = C_2(t_2')$.

Now we need to prove that $\forall l \in \text{view}_{H_2'}(\Delta(Sc_2') \ \cup \ \Delta(Sc_2)) \ f'(H_1') \cong_l H_2'$. This is equivalent to proving

L$\forall l \in \text{view}_{H_2'}(\mathcal{L}_{C_2} \ \cup \ \Delta(S_2') \ \cup \ \Delta(S_2)) \ f'(H_1') \cong_l H_2'$

Consider a heap address $l \in \text{view}_{H_2'}(\mathcal{L}_{C_2} \cup \Delta(S_2) \cup \Delta(S_2'))$. Proving the above statement is equivalent to showing that for all such l, $f'(H_1')(l) = H_2'(l)$.

- If $l \in \text{view}_{H_2'}(\Delta(S_2) \cup \Delta(S_2'))$ then using (g) we get $f'(H_1') \cong_l H_2'$.
- Else $l \notin \text{view}_{H_2'}(\Delta(S_2) \ \cup \ \Delta(S_2'))$ then by Lemma 2 we know that $H_2(l) = H_2'(l)$. For the theorem we only need to consider the case when $l \in dom(H_2)$. Since $\text{view}_{H_2'}(\Delta(S_2) \ \cup \ \Delta(S_2')) = \text{view}_{f'(H_1')}(\Delta(f(S_1)) \ \cup \ \Delta(f'(S_1'))) = \text{view}_{f'(H_1')}(\Delta(f'(S_1)) \ \cup \ \Delta(f'(S_1'))$ (By applying (f)).
  Therefore $l \notin \text{view}_{f'(H_1')}(\Delta(f'(S_1)) \ \cup \ \Delta(f'(S_1')))$. Thus applying Lemma 2 (to the transition $f'(S_1) \to f'(S_1')$) we get $f'(H_1)(l) = f'(H_1')(l)$. From (b) and (f) we can show that $f'(H_1)(l) = f(H_1)(l) = H_2(l)$. Thus we get $f'(H_1')(l) = H_2'(l)$.

Thus we have shown the theorem for the inductive case. $\square$